The Claim Is the Promise

How great claims build trust, loyalty and long-term growth

Ian Hughes

2026

The Claim Is the Promise

How great claims build trust, loyalty and long-term growth

First edition, 2026

The evidence in this book is drawn from six years of independent claimant tracking, covering more than 150,000 motor and home customers. Every figure is subject to the notes in Appendix A.

hello@claims.company

Contents

Introduction: The promise nobody reads

I have spent my working life building promises.

That is what marketing is, once you strip away the noise. A brand is a promise. It tells a customer what to expect, how they will be treated, and who they are dealing with. I have spent years learning how to make those promises land. How to make them feel true before they have been tested. In most industries, that is enough. The promise and the product arrive close together. You buy the coffee, you drink the coffee, and the brand is either honest or it is not, there and then.

Insurance is not like that.

In insurance, the promise and the moment of truth can be decades apart. You buy a policy today. You pay for it every year. And you hope, quietly, that you will never need it. The thing you have actually bought is a sentence. If something goes wrong, we will be there. Everything else, the adverts, the logo, the friendly voice on the phone, is written around that single sentence. And here is the paradox that has stayed with me for most of my career. The customer is handed over at the exact moment the promise is finally tested. The claim. Brand can reach almost everywhere. It cannot reach the one place that matters most.

That bothered me. As a marketer, I could shape everything up to the

claim, and almost nothing at the claim itself. So I started paying attention to what happens there. And the more I looked, the more I realised that claims are not really about claims at all. They are about trust.

Insurance is a contract of trust in the purest form I know. You pay now for a promise redeemed later, and perhaps never. There is no other everyday purchase quite like it. You cannot inspect the thing you are buying. You cannot try it before you commit. You simply have to believe that when the worst happens, the promise will hold. The claim is where that belief is tested. It is the redemption of the contract. It is the moment the customer finds out whether the promise was real.

I am writing this book at a moment when that trust is under strain. The industry is increasingly accused of breaking faith with the people it serves. Of taking the premium and resisting the claim. Of serving itself first. I am not going to pretend that criticism has no cause, and I am not going to spend these pages defending the industry against it. Defence is not leadership. The far more useful question, the one that has fascinated me, is simpler and harder. What does trust actually need? How does it work? And if we understood that properly, could we build more of it, on purpose?

To answer that, I needed something better than instinct. I needed a way to take trust apart and see how it is made.

Years ago I came across a formula that has shaped how I think ever since. It comes from a book called *The Trusted Advisor*, by David Maister, Charles Green and Robert Galford. They reduced trust to an equation. Trust is credibility plus reliability plus intimacy, all divided by self-orientation. Four simple terms. Do you know what you are doing. Do you do what you said. Is it safe to be open with you. And, beneath all of it, are you in this for me, or for yourself.

I have carried that equation into boardrooms, into brand strategies, into every serious conversation I have had about why people believe in one company and not another. And when I finally held it up against the insurance claim, it fitted almost perfectly. A claim tests credibility,

because the customer finds out whether you truly know your craft. It tests reliability, because you either do what you promised, when you promised, or you do not. It tests intimacy, because a claim is a moment of real vulnerability, and the customer learns whether they are in safe hands or merely being processed. And it tests self-orientation most of all, because every accusation ever levelled at this industry is, at heart, a suspicion about the denominator. The fear that when it mattered, the insurer was serving itself.

That is the lens this book will use, from the first chapter to the last.

There is one more thing I want to say before we begin, because it runs underneath everything that follows. Claims are not processes. They are people. On one side, a person at a moment of loss, often frightened, often out of their depth, whose feelings are as much a part of the claim as the paperwork. On the other side, the people who deliver the experience, whose care or indifference the customer feels immediately and remembers for years. It has become fashionable to talk as though technology will soon handle all of this. It will handle a great deal, and I will make the case later that it should. But I want to plant a flag here at the start. In a world where machines can run the process, the human moments become more valuable, not less. They become the scarce thing. And scarce things are where the promise now lives.

So this is not really a book about claims, though claims are on every page. It is a book about trust, told through the one moment where trust is truly won or lost. It is written by a marketer who spent a career on the promise and became fascinated by the place the promise gets tested. And for the first time, it is a book that does not have to rely on my opinion alone. Because we can now see claims experience across the whole market, at a level of detail that has never existed before, and let the evidence speak.

Let me tell you where that evidence comes from, because you are entitled to know what I am standing on. For more than six years, a study called the Insurance Behaviour Tracker has been asking ordinary people about their insurance. Not what the industry assumes they feel, but what they

themselves report, in their own words, while the experience is still fresh. It is run by Consumer Intelligence, the business I led. In that time it has heard from more than 150,000 motor and home customers, over 30,000 of them claimants describing a real claim they had just lived through. To my knowledge, that is one of the largest bodies of claimant experience ever gathered in this country. And it has a quiet feature that matters more than it might sound. It is not a claims survey. It asks people about their insurance in the round, and comes to their claims experience along the way. That means the people who answer are not only the delighted and the furious, the ones with a story they are burning to tell. They are the whole spread, including the great, quiet middle whose experience was simply fine, and whom a claims survey almost never hears from. As we will see, that middle is where a surprising amount of loyalty is quietly won and lost. And the tracker has one more feature, on which much of this book turns. It follows claimants forward to their next renewal. So what a person experienced at the claim can be tied to what they actually did next, at the same price. Not what they told a researcher they might do. What they did. So when I make a claim about claims in the pages that follow, that is what sits beneath it. Not my instinct. Their experience, counted at a scale no single insurer could ever reach.

And what that evidence shows is worth setting down before we begin, because it is the spine of the book. Follow claimants forward and the distance between a good claim and a poor one is somewhere between twenty-six and forty-plus percentage points of a customer's future. Points of staying or leaving, not points of sentiment. And the gap survives the obvious objection. It is not a price effect in disguise. Compare customers who saw the same change in price, of similar age, in the same part of the country, and it barely moves. There is a dividend sitting inside the claim. Most of the industry cannot see it, because the instruments it watches were never built to look.

The book travels in four parts.

The first makes the argument. Why the claim is the only test that counts.

And why, in a market where nearly every claim is paid, paying has become the price of entry, and how it felt has become the thing that separates one insurer from another.

The second lays out the dividend. What a good claim is worth in motor and in home, told separately, because they are not the same story. Motor holds a surprise: a well-handled claim can leave a customer more loyal than never claiming at all, and the effect is strongest exactly when the renewal price rise is biggest. Trust, we will find, is banked at the claim and spent at the price rise. Home holds a harder truth: even a good home claim leaves a bruise, and the prize there is halving the damage. Between them sits that quiet middle, the claimants who will tell you everything was fine, and who leave at close to the rate of the angry. Fine, it turns out, is not fine. The part ends with the arithmetic. A method any finance director can rerun with their own numbers, and the question it forces: if the dividend is this large, why has almost nobody seen it?

The third part goes to where the dividend is actually won and lost. Into the journey itself: the first contact, the silences, the moment of being seen as a person rather than a case number. Into measurement that deserves a board's trust. And into the place where process meets people, because the answer lives in neither alone.

The fourth faces what is arriving. The cost pressure squeezing the promise, told from the industry's side of the ledger and told fairly, because the case for handling claims well is only credible from someone who understands what handling claims costs. The machine, because AI is coming to claims whether we are ready or not, and it will either industrialise care or industrialise its absence. The regulator, whose measures count what firms did but cannot yet see how it felt, or what the customer did next. And the investor, for whom claims quality turns out to be a leading indicator of retention economics that statutory reporting has never carried.

That widening cast is deliberate. I wrote this book for the people who run claims, and for the people who run the companies that contain them. But

a good yardstick can be read by a customer, a leader, an investor, and in time perhaps a regulator, and the later chapters serve each of them in turn. One promise on the numbers themselves: I have kept them deliberately spare in the text, and every figure that appears in these pages sits at the back of the book with its source, its base and its caveats. Where the data cannot yet answer a question, the closing pages say so, and say what we will know next.

The claim, it turns out, is the promise. Let me show you what that means, and what a great one is worth.

Chapter 1: The only test that counts

Think about the last time a company truly let you down. Not a minor irritation, but a real one. A moment when you needed them to be what they had promised to be, and they were not.

You can probably still feel it. The name of the company, the tone of the person you dealt with, the exact point at which you realised you were on your own. We forget a thousand smooth transactions. We remember the one that failed us, and we remember it for years. We tell other people about it. We change who we give our money to because of it.

Psychologists have a name for this lopsidedness, and it is worth knowing, because the whole economics of claims rests on it. Daniel Kahneman, who won a Nobel prize for understanding how people actually think, showed that we do not remember experiences evenly, minute by minute. We remember two things: the most intense moment, and how it ended. He called it the peak-end rule. Now look at an insurance claim through that lens. The claim is the most intense moment the customer will ever have with you. And for too many customers, it is also the end. The one event is both the peak and the finish, which means the claim does not just influence the memory of the relationship. For most customers, the claim *is* the memory. Everything else, the years of renewals, the direct debits, the documents filed unread, dissolves into background. Other

researchers have put the same truth more bluntly: bad is stronger than good. It takes a number of good experiences to outweigh a single bad one, because we are built to remember the snake and not the path. For an insurer, this is brutal arithmetic. A decade of quiet, flawless renewals can be spent in a fortnight of unreturned phone calls.

Now think about how much of the insurance industry's effort is aimed at the moment before any of that can happen.

We spend enormous sums making the promise attractive. The advertising, the sponsorships, the price comparison rankings, the clever design of the quote journey. All of it is pointed at a single event, the sale. Getting the customer to say yes. And it works. The industry is very good at winning the quote. But the quote is not where reputations are made. The quote is where a relationship begins. It tells you almost nothing about whether that relationship will be honoured.

Marketers of my generation were raised on a story from the airline industry. Jan Carlzon, who ran Scandinavian Airlines in the 1980s, worked out that his company was created in what he called moments of truth: tens of millions of brief encounters a year, each one a few seconds long, in which a passenger met a member of his staff and formed a judgement about the whole airline. His insight changed service businesses everywhere. But insurance is a stranger case than Carlzon's, and I think a more extreme one. An airline has millions of moments of truth. An insurer has, in effect, one. In any given year, fewer than one customer in ten will make a claim; on our own tracker's count, it is nearer one in sixteen. Most customers will go years, sometimes decades, without ever testing the thing they are paying for. Two consequences follow, and they shape this entire book. First, when the moment finally arrives, it carries the accumulated weight of every untested year. The customer is not just asking whether you will fix the car. They are finding out, all at once, whether the fifteen years of premiums meant anything. Second, the claimant becomes the witness for everybody else. The fifteen in sixteen who did not claim this year cannot judge you on their

own experience, so they judge you on the experience of the one who did. The story told at the school gate, in the office kitchen, at Sunday lunch. The claim is the industry's only moment of truth, and it is performed in front of an audience.

The honouring happens at the claim.

I have come to believe that the claim is the only test that truly counts, and I want to be precise about why. Everything before the claim is a promise. Everything after it is a memory. The claim is the only moment where the promise becomes real, in front of the customer, when it matters to them most. It is the hinge on which the entire relationship turns. And unlike the quote, it is not a moment the customer chooses lightly or forgets easily. They come to you because something has gone wrong in their life. A crash. A flood. A theft. A loss. They are not shopping. They are hoping.

It is worth pausing on what that day actually looks like, because everything in this book happens inside it. Nobody rings an insurer's claims line from a sun lounger. They ring from the hard shoulder with the hazard lights ticking, copying a stranger's registration number into their phone in the rain. They ring because a son who passed his test three weeks ago has just called, voice shaking, saying Dad, I'm so sorry. They ring standing in a hallway that smells of wet plaster, or outside a back door that no longer locks because someone forced it. The customer who reaches you is having one of the worst days of their year, occasionally one of the worst of their life, and they are about to conduct the only inspection of your product they will ever carry out. In that state. That is the examination room. No brand campaign gets to sit in it. Only the claim does.

That is what makes the claim so different from every other point of contact, and so much more revealing. At the quote, the customer is comparing prices. At the claim, the customer is finding out who you really are.

Here is where the trust equation earns its place, because a claim tests all four of its terms at once, and it tests them under pressure.

It tests credibility. When a customer makes a claim, they discover, often for the first time, whether the company on the other end actually knows what it is doing. Whether the person handling their case understands the policy, the process, the options. Whether the promises in the brochure were backed by real competence, or were just words. A customer can forgive a great deal, but they can tell within minutes whether they are in capable hands.

It tests reliability. A claim is a sequence of small promises. We will call you back. The assessor will come on Tuesday. Your car will be ready by Friday. Each one is a chance to do what you said, when you said it. Or not. Reliability is not a grand gesture. It is the quiet accumulation of kept word after kept word, and the customer is counting, even when they do not realise they are counting.

It tests intimacy, which in this setting means something very human. Can I be vulnerable with you, and be safe. A person making a claim is often frightened, sometimes ashamed, usually out of their depth. They are not at their best. Whether they are met with warmth or with indifference, whether they feel handled or feel helped, is the part of a claim they will describe most vividly, long after the money has been paid. This is the term the industry measures least, and it may be the one that matters most.

And beneath all three, the claim tests self-orientation. This is the denominator, and it is the one that decides everything. The customer is asking, silently, a single question. When it came to it, were you on my side, or your own? Every good claim answers that question one way. Every bad one answers it the other. And because insurance has spent years accused of serving itself first, this is the term where the whole industry's reputation is being contested, one claim at a time.

Notice, too, what the shape of the equation implies, because it is not a detail. Credibility, reliability and intimacy sit on top. They add. You can build them up term by term, and a weakness in one can be partly carried by strength in another. Self-orientation sits underneath. It divides. A single moment in which the customer decides you were

serving yourself does not subtract a little from the trust you have built. It divides the whole of it. This is why a claim can be technically perfect and still destroy the relationship. The car was repaired. The payment arrived. Every box on the process map was ticked. But somewhere in the middle the customer caught the company, or believed they caught it, putting its own interests first. A settlement that felt engineered downwards. A delay that felt designed. From that point the competence stops counting in the customer's favour and starts counting against you, because a capable company that is not on your side is worse than an incompetent one. The denominator is where claims are truly lost. Most of the industry's measures never look at it.

Four terms, all tested together, in the single moment the customer will never forget. That is why the claim is not a back-office function. It is the front line of trust. It is, quite literally, where the brand is proven or exposed.

I understand why the industry has not always seen it this way. Claims sit in operations, measured in cost and cycle time and leakage. Marketing sits somewhere else entirely, measured in awareness and acquisition. The two halves of the business rarely meet. And so the most powerful brand moment an insurer will ever have, the claim, is managed as though it were a cost to be contained rather than a promise to be kept. That separation is one of the quiet tragedies of this industry, and closing it is one of the things this book is for.

There is a structural reason the divide persists, and it is worth naming, because it is nobody's fault and everybody's problem. Claims are the largest cost in an insurer's accounts. The greater part of every pound of premium collected goes back out of the door through the claims line. Anything that large attracts, quite properly, the disciplines of cost: budgets, targets, efficiency programmes, leakage reviews. And so the people who run claims are asked, year after year, to make the number smaller, and are promoted when they do. Meanwhile the value a claim creates, the loyalty, the advocacy, the story told at Sunday lunch, lands

somewhere else entirely, in renewal rates and acquisition costs that are counted in different meetings by different people who do not think of themselves as being in the claims business at all. The cost of a claim is visible in one place. The value of a claim is scattered everywhere else. Nobody decided to hide the moment of truth. The organisation chart did it on everyone's behalf. And a thing hidden by structure needs more than goodwill to reveal it. It needs measurement, which is where this book is eventually heading.

None of this is a criticism of the people doing the work. Most claims professionals I have met understand the weight of the moment better than anyone. They feel it every day. What they have lacked is not care. It is evidence. A way to show the board that the claim is not merely an expense, but the single greatest opportunity the business has to build the trust that keeps customers and drives growth. For a long time, that evidence did not exist. The claim was the most important moment in the relationship, and also the least visible.

That is the part that has changed, and it is why I felt able to write this book now rather than ten years ago. For the first time, we can see claims experience across the whole market. Not one insurer's internal numbers, defended and disputed, but the market itself, at a level of detail that has never been available before. We can see how widely experience varies from one insurer to the next. We can see what a good claim does to loyalty, and what a poor one costs. We can, at last, take something that leaders have always sensed in their gut and put it on the table as fact.

So this is the argument the rest of the book will build. The claim is the only test that counts. It is where the promise is kept or broken, where all four terms of trust are decided at once, and where the future of the relationship is written. Treat it as a cost, and you will manage it down. Treat it as the moment of trust it truly is, and you will build something a competitor cannot buy with a bigger advertising budget.

Before we get to what a great claim is worth, though, there is a myth we need to clear out of the way. Many people in this industry believe the

real test of a claim is whether you pay it. It is not. And the evidence for that is the most surprising place to start.

Chapter 2: Everybody pays, so the prize is elsewhere

Ask most people in this industry what makes a good claim, and sooner or later the answer comes down to one thing. You pay it. The claim is a promise to pay, so a good claim is one where the money arrives. A bad claim is one where it does not. Simple.

It is also, on the evidence, almost beside the point. This chapter explains why, because it is the most useful ground clearing we can do. Once you see it, the whole competitive picture changes shape.

Here is the finding, and it surprised me when I first sat with it. Claims are, overwhelmingly, paid. Across the market, the vast majority of motor claims that are made are accepted. And more importantly, the firms are bunched tightly together at the top of that scale. It is not that one insurer pays nine claims in ten while another pays six. They nearly all pay at rates so high, and so close to one another, that acceptance cannot be the thing that separates a great insurer from a poor one. On this measure, the whole industry looks the same.

Let me show you where that finding comes from, because the source matters. Since the regulator began publishing what it calls Value

Measures, anyone has been able to look up, firm by firm, the things a firm can count about its own claims: what share of claims it accepts, how often its customers claim, what the average payout is, and how many complaints it receives. It is the industry's official ledger, and it is where the acceptance picture above comes from. What we have been able to do, for the first time, is take that ledger and lay it alongside the tracker I described in the introduction: the same firms, matched name by name, with the regulator's record of what each firm *did* sitting next to its own customers' account of how it *felt*. Eleven motor firms and eight in home, matched at the level of the underwriter the regulator reports, read both ways at once. To my knowledge nobody has been able to put those two pictures side by side across the market before. And when you do, something jumps off the page.

Picture the two measures side by side. Line the matched motor firms up by their acceptance rate, the regulator's measure of whether claims get paid, and they sit almost on top of one another, huddled in a narrow band near the top of the scale, so close together that you would struggle to tell one from another with the naked eye. Now line the very same firms up by their claimant satisfaction, their own customers scoring the experience of the claim, and they fly apart, scattered from the low seventies to the high eighties. Between the best firm and the worst, the spread runs to nearly twenty points. Nearly twenty points, between companies whose acceptance rates are indistinguishable. Two readings of the same firms, and two completely different truths. On what firms did, the market is a monolith. On how it felt, the market is a landscape, with summits and valleys a customer can feel from inside a single phone call. If there is one contrast to carry into your next board meeting, make it that one. Paying is the price of entry. Feeling is the differentiator.

Sit with that for a moment, because it undoes a comfortable assumption. If almost everyone pays, and everyone pays at roughly the same rate, then paying cannot be your differentiator. It is the price of entry. It is the thing you have to do to be allowed to compete at all. A firm that congratulates itself on its acceptance rate is congratulating itself on turning up.

It is worth asking why everyone pays, because the answer is not sentimental. It is regulatory. Years of rules on fair value and, more recently, the Consumer Duty have made wrongly declining a claim one of the most dangerous things an insurer can do. The floor is held up by supervision, not virtue. And when a corner of the market has failed the paying test, the regulator has moved. Guaranteed asset protection cover, the add-on that pays out the gap on a written-off car, was paying out so little of what it collected that the regulator intervened, and payout rates rose sharply almost immediately. That is the fate of any product that fails on acceptance: it becomes a regulatory project. Mainstream motor and home passed that test long ago. Which is precisely why the contest has moved. An airline does not win passengers by advertising that its planes land. Landing is assumed. What is left to compete on is everything the passenger actually remembers.

I should add one note of caution here, because it matters. Even the published acceptance figures come with a warning attached. The regulator itself has flagged that some of its acceptance data should be read with care, because firms do not all count and categorise claims in the same way. So the true picture may be even flatter than it looks. But the direction of travel is not in doubt. Acceptance is high, it is common, and it does not tell customers apart in any way they can feel.

I should also be candid about the boundaries of this finding, because rigour matters more than a tidy headline. It is cleanest, and strongest, in motor. Motor is a well defined product. What is covered, what is not, and what a claim looks like are relatively clear, which is why acceptance clusters so tightly. Home and travel are different animals. Their products vary far more from one policy to the next. Their payout rates are more spread. The way claims are counted is less consistent. And home, in particular, is a more emotive product, bound up with the place a person lives rather than the car they drive. So when I say that almost everyone pays, I mean it most confidently of motor, and it is why the contrast I have just described holds for motor: in home, the regulator's own acceptance measure really does vary from firm to firm, so the flat line is not there.

The wider principle, that experience matters more than acceptance, holds across all three. But the acceptance data itself is at its flattest, and its most reliable, in motor. That is where this book plants its feet, and where its evidence is strongest. Home, I promise, is not being ducked. It gets a chapter of its own later in the book, and a harder truth of its own too, because the home story turns out to be different, and not in the direction you might guess.

Now let me connect this to trust, because this is where it becomes interesting rather than merely surprising.

Remember the denominator of the trust equation. Self-orientation. The customer's silent question at the claim. Were you on my side, or your own? Paying the claim is the industry's most basic answer to that question. It says, at the minimum, we did not serve ourselves at your expense. We honoured the contract. And that matters. A refused claim, rightly or wrongly, lands as a betrayal. And the cost of getting the rest of it wrong is not vague or sentimental either: a poor claim roughly doubles the rate at which customers walk away, in motor and in home alike. Not a few points. Double. We will take that number apart properly in the next chapter, because it deserves its own room. For now, hold the shape of it: the difference between a good claim and a poor one is not a rounding error in a retention report. It is one of the largest levers in the business.

But here is the point I most want you to take from this chapter. Paying the claim clears the floor. It does not win the room. Answering the self-orientation question with a bare yes, we paid, is the beginning of trust, not the achievement of it. The customer notices that you paid. Then they notice everything else. How long it took. Whether they had to chase. Whether they were treated as a person or a case number. Whether anyone seemed to care that their week, or their year, had been upended. The payment settles the contract. The experience settles the relationship.

This is why two customers can both have their claims paid, in full, and walk away feeling completely differently about their insurer. One felt

held. One felt processed. Both got the money. Only one will still be a customer in three years, and only one will tell their friends they were looked after. The cheque was identical. The trust was not.

I find this liberating, and I hope you will too. Because if paying were the whole game, there would be very little left to compete on. Everyone pays, so everyone would be equal, and claims would be exactly the cost centre the industry has long treated it as. But paying is not the whole game. It is barely the start of it. Which means the real prize, the thing that actually decides whether a customer stays and recommends and comes back, is still sitting there almost entirely unclaimed. It is not price or acceptance. It is the experience of the claim itself, and how it makes people feel.

That prize is bigger than price, and it is more durable. A competitor can undercut your premium tomorrow. A competitor can outspend your advertising next quarter. What a competitor cannot easily copy is a claims experience that customers trust, because that is built from a thousand human moments and the leadership that shapes them. It is the one advantage in this industry that does not evaporate the moment someone else lowers their price.

So we can put the myth to rest. The test of a claim is not whether you pay it. Almost everyone passes that test. The test is how you pay it, and who the customer discovers you to be while you do. That is where insurers differ, enormously, and it is where the rest of this book will live.

Which raises the obvious question, and it is the one worth real money. If experience is the prize, what is a good claims experience actually worth? Not in sentiment, but in customers kept and growth earned. For the first time, we can answer that with numbers. And the answer is larger than most boards would dare to guess.

Chapter 3: What a good claim is worth

The last chapter left a question hanging. If the experience of a claim is the real prize, what is it actually worth? Not in warm words, but in the things a board cares about. Customers kept. Growth earned. Let me answer it as directly as the data allows.

Start with how much claims experience varies, because if it did not vary there would be nothing to compete for. At the level of the whole market, satisfaction with claims is reasonably high. Somewhere around eight in ten claimants come away satisfied. That is the number an industry tells itself to feel reassured. But the market average hides the thing that matters. Underneath it, the range is wide. At the insurers customers rate most highly, close to nine in ten claimants come away satisfied. At those they rate least, it is nearer seven in ten. Almost twenty percentage points of satisfaction separate the best from the worst. Same product. Same promise. A completely different experience.

Twenty points is not a rounding error. It is the distance between a claimant who feels looked after and one who feels let down, and it is almost entirely within a firm's control. The spread is the proof that experience is a choice, not a fixed cost of doing business. Some firms have chosen well. Others have not yet realised there was a choice to make.

Now to the number itself. The one worth taking to your board. When you follow claimants forward to their next renewal and look at what they actually do, the pattern is stark. Across two years of renewals, roughly three in ten motor customers who had a good claim left their insurer. Among those who had a poor claim, it was nearly six in ten. Home tells the same story. A poor claim roughly doubles the rate at which customers walk away. Put as a gap, that is twenty-nine points of loyalty in motor and twenty-six in home. And when you measure it another way, comparing claimants who came away satisfied with those who came away dissatisfied, the renewal gap stretches beyond forty points. However you cut it, the distance between a good claim and a poor one is somewhere between twenty-six and forty-plus percentage points of a customer's future. Set that against everything else a business does to hold on to its customers. The loyalty schemes. The retention discounts. The win-back campaigns. Few of them shift loyalty by anything close.

I can hear the objection, because I have heard it in every room where I have laid these numbers out. Perhaps it is not the claim at all. Perhaps the customers who left were just the ones whose price went up at renewal. It is a fair challenge, and it deserved a proper answer, so we tested it. Compare like with like, customers who saw the same change in price, of similar age, in the same part of the country, and the gap does not disappear. It barely moves. Around thirty points in motor. Around twenty-five in home. The dividend is a claims effect, not a price artefact. The promise, kept or broken, moves the customer all on its own.

Let me be precise about those two measurements, because a careful reader deserves to see the joins. They are different rulers laid against different groups. The first follows customers whose claims were good or poor and counts who actually switched: that is the conservative number, the one that survives the price test, and I will treat it as the floor. The second compares the satisfied with the dissatisfied and counts who renewed: a wider net, and a wider gap, and I will treat it as the outer bound. I will never add them together or swap one for the other, and if you remember only one, take the floor. Everything this book goes on to

argue stands comfortably on the smaller number.

There is a question the numbers alone cannot settle, which is what the customers themselves would say moved them. So we asked. When customers who shopped around are asked why, most reasons are the familiar ones. Price went up. Time for a change. But three of the reasons a customer can give point straight back at the claim, in their own words: the claims service was poor. The cover was not what I expected when I claimed. My claim was rejected. And here is what the answers show. Having made a claim at all, good or bad, roughly doubles to triples the chance that a shopper gives one of those claims-driven reasons. In motor, about one in three shoppers who had claimed cited one, against fewer than one in five who had never claimed. In home, more than two in five, against roughly one in seven. The claim, once lived through, stays on the shopping list. But I owe you the finding's other half, because it is just as instructive. Customers whose claims went badly give these reasons no more often than customers whose claims went well. The words record that the claim happened. They do not record how it was handled. For the verdict on handling, you have to watch the feet. The behavioural gap, the one that survives the price test, is written by the quality of the claim; the customers' own words confirm only that the claim was the event that set them shopping. So the two kinds of evidence divide the labour, and that is the true shape of the case. What customers say puts the claim in the dock. What customers do delivers the verdict on how it was handled. We will meet this pattern again, because it runs through the whole book: people are poor narrators of their own loyalty, and behaviour is the better witness.

I am careful not to leave this in the language of satisfaction, because satisfaction sounds soft, and boards are right to be wary of soft. So let me put it in the language of growth. A retained customer is the cheapest growth there is. You have already paid to win them. Keeping them costs a fraction of finding their replacement. A book of business that holds on to its satisfied claimants compounds, quietly, year after year. A book that loses them has to run just to stand still, buying new customers to replace

the ones a poor claim drove away. The claim is a growth engine, not a cost line, and much of the industry has been running it in reverse.

I will not put a single pound figure on it here, and I will tell you why. The moment you turn a retention gap into a precise sum of money, you have to make assumptions about your own book, your own claim rates, your own margins. Those are your numbers, not mine, and the honest thing is to hand you the method rather than a headline. Take the gap. Apply it to the size of your book and the share of your customers who claim in a year. You will have a figure that is defensible precisely because it is yours. What I can tell you, with confidence, is that for any insurer of scale that figure is large enough to change how you think about claims. And I will not leave the method as a sketch. Later in the book, once the whole dividend is on the table, there is a chapter that walks through the arithmetic properly, line by line, with worked examples you can hold against your own accounts.

If you ask which part of the claim does the most to earn that loyalty, I spent years assuming the data would point to one term of the trust equation. Reliability. Doing what you said you would do, when you said you would do it.

Reliability is not glamorous. It does not show up in an advert. But it is the term the customer can feel at every single step. The call that came when it was promised. The update that arrived before they had to ask. The timescale that held. Each kept promise is a small deposit of trust, and by the end of a well run claim the customer has watched you keep your word a dozen times over. I held that belief provisionally through the early drafts of this book, and said so, because I intended to test it properly, statement by statement, against what actually separates the customers who stayed from the ones who left. We have now run that test, and the answer changed my mind in the best possible way. Reliability does not win. Neither does anything else. The kept promise and the human warmth of the claim, the sympathy, the friendliness, the sense of being in kind hands, separate the stayers from the leavers by almost exactly the

same amount. No honest reading of the data crowns one over the other. And I find that more useful than the tidy answer I expected, because it closes a door the industry has been tempted by for years: the idea that a claim could be won by process alone, with the timescales kept and the humanity engineered out. The data says no. The warmth carries as much of a customer's future as the kept word. Later in the book we will walk through that test statement by statement, and you will see the whole ranking.

The same mechanism runs in reverse, and it is worth naming, because it sets the stakes. A claim full of broken promises does not merely disappoint. It pushes a customer away, and it does something worse. It gives them a story to tell. People who have a poor claim rarely keep it to themselves. They warn their friends. They post their reviews. They become a small engine of lost reputation, working against every pound the marketing team spends to build it. And this, too, no longer has to rest on assertion. The survey asks customers how likely they are to recommend their insurer, which means advocacy can be read by claim group: the well-handled next to the badly-handled, the enthusiasts next to the warners. The reading is stark. A good claim creates advocates: among good-claim customers, net recommendation runs strongly positive in both products. A poor claim creates the opposite, and the swing between the two is enormous. Around eighty points of net recommendation in motor. More than a hundred in home. So the gap cuts both ways. Handled well, the claim is your best growth event. Handled badly, it is your most expensive one.

I do not raise the downside to alarm you. I raise it because it is the clearest possible case for treating claims as a leadership priority rather than an operational afterthought. The upside and the downside are both large, both real, and both sitting inside a moment most firms currently manage for cost. That is not a problem to be feared. It is an opportunity that has been hiding in plain sight, waiting for someone to lead it.

And yet the number I have given you in this chapter, large as it is, is

still not the most surprising thing the data has to say. So far we have compared good claims with poor ones. The stranger result appears when you compare claimants with the customers who never claimed at all, and then watch what happens at the worst possible moment: the renewal where the price goes up. What a good claim does in that moment is something I did not expect, and it changes how you should think about every pound you spend on claims. That is the business of the next chapter.

Chapter 4: Trust banked, trust spent

Every insurance relationship contains one letter the customer actually reads.

It is not the policy. We established in the introduction that nobody reads the policy. It is the renewal notice, and they read it for one reason: it contains a number, and the number has usually gone up. In that moment, every insurer in the market is reduced to a single question in the customer's mind. What am I actually paying for? A comparison site is thirty seconds away. The whole industry has trained customers, for two decades, to answer a price rise by shopping. The renewal letter with a bigger number on it is the single most dangerous moment in the relationship. More customers are lost there than anywhere else.

So here is the experiment worth running, and thanks to the tracker following claimants to renewal, we could run it. Take customers facing that dangerous moment. Hold the moment constant: the same change in price, like for like. Then divide them by one thing only, their claims history. Some had never claimed at all. Some had claimed, and it went well. Some had claimed, and it went badly. Now watch who stays.

You might expect the never-claimed to be the loyal ones. They have had no reason for grievance. Their relationship is unblemished, all quiet

direct debits and renewal letters. The claimants, by contrast, have been through something. Even a good claim is an interruption, an intrusion of paperwork and phone calls into a relationship designed to be silent. Common sense says the customer with the unblemished record stays, and the one who has been through the mill looks around.

In motor, common sense is wrong, and it is wrong in the most interesting way I have found in this whole body of data.

There is one more finding here, and I offer it carefully, because it surprised me, and because it comes with an asymmetry that I am not going to smooth over.

In motor, a well-handled claim can leave a customer more loyal than one who never claimed at all. At the same price rise, customers who had a good claim switched less than customers who had never tested the promise. And the effect was strongest exactly where loyalty is hardest to hold, among customers facing the largest price increases, where good claimants were some twenty points less likely to leave than the never-claimed. Think about what that means. The claim is the one moment an insurer can prove the promise is real. And a proven promise carries a customer through the very renewal that would otherwise have lost them. Trust is banked at the claim, and spent at the price rise.

Look at the shape of the finding, because the shape is the argument. When the renewal price barely moves, the good claim buys you almost nothing you can measure. A few points at most, and statistically you could not swear to them. When the rise is modest, still nothing you could take to a board. But when the rise is large, more than pocket change, the kind of increase that sends a customer straight to a comparison site, the gap opens to some twenty points. The effect is not spread evenly across the relationship. It concentrates itself, precisely, at the moment of maximum danger. A good claim is not a general feeling of warmth. It is a specific asset that gets drawn on at a specific moment, and that moment is the renewal that would otherwise have lost the customer.

Why would it work this way? Go back to the customer’s question at the

renewal letter. What am I actually paying for? For the never-claimed customer, that question has no answer. They have paid for years and received, as far as they can see, nothing but paperwork. The promise is still just a promise, untested, weightless against a saving of eighty pounds on a comparison site. But the good claimant has an answer on file. They know exactly what they are paying for, because they have seen it. The hire car that arrived. The voice that took charge on a bad morning. The thing that went wrong and then, remarkably, got put right. When the price rise poses its question, the never-claimed customer hears silence, and the good claimant hears a memory. That is the entire difference, and it is worth twenty points of loyalty at the worst moment in the relationship.

This is also, I think, the deepest meaning of the word dividend. A dividend is what an asset pays out, later, to whoever held it. The claim is where the asset is created. The renewal is where it pays. And like any asset, it exists whether or not anyone puts it on a balance sheet. Most insurers are holding this asset right now, unmeasured, unmanaged, and unpriced.

The timing of this finding matters too, because the ground has moved under retention. For years, much of the industry held on to customers the engineered way. New customers got the attractive price, loyal ones paid more each year, and inertia did the rest. The regulator ended that in January 2022, when the rules on pricing practices banned the loyalty penalty outright. Whatever you think of that intervention, its consequence is plain: retention can no longer be engineered through the back book. It has to be earned. And the moment you have to earn loyalty rather than harvest it, the question becomes what actually earns it. This chapter's answer is uncomfortable for anyone whose retention strategy lives entirely in the pricing department: the biggest earning event available to an insurer is not a discount. It is a claim, handled well, sitting in the customer's memory when the difficult renewal arrives.

Now I need to close a door, firmly, because I can already hear it creaking open. There is a cynical reading of this finding, and it goes like this:

good claims make customers tolerant of price rises, therefore invest in claims and raise prices with impunity. Let me say as plainly as I can that this is not the lesson, and a firm that acts on it will destroy the very asset this chapter describes. Remember the denominator. The entire value of a good claim is that it answers the customer's oldest suspicion, that the insurer serves itself first. Use the trust a claim has banked as cover for extracting price, and you have converted the proof of your good faith into an instrument of the very self-orientation it disproved. Customers are not stupid. Markets are not stupid. And the regulator, whose fair-value rules exist for exactly this, is not stupid either. The finding says a kept promise carries customers through fair prices honestly arrived at. It does not say, and I will not let it be read to say, that a kept promise is a licence to charge more. Trust banked and then spent on the customer's behalf compounds. Trust banked and then spent against them is embezzlement, and it gets discovered.

What the finding does license is something more valuable and entirely respectable. It licenses the claims director and the pricing director to stop pretending they run separate businesses. In most insurers I know, claims and pricing sit in different buildings, report through different lines, and meet mainly at budget time, where claims is a cost to be argued down. The data says they are one system. The pricing team decides the size of the question the renewal letter asks. The claims team decides whether the customer has an answer. Neither controls retention alone, because retention is made in the space between them. An insurer that prices renewals with no idea which customers are carrying a banked claim, and runs claims with no idea what those customers will face at renewal, is playing both halves of the game blind. The first firm to run them as one system will see what the market cannot: exactly where its trust is banked, and exactly when it will be drawn on.

The limits of this finding need flagging before we move on, because this book does not do conjuring tricks. The price change we hold constant is the one the customer remembers, and memory is not a ledger. Customers who had a bad claim may also remember their price rise as bigger than it

was; controlling for remembered price reduces that distortion but cannot remove it entirely. The strongest reads come from claims followed over two years of renewals; the sharpest price-band cuts rest on a single year. And every figure here is motor. I have not shown you home yet, and that is deliberate.

Because home is different, and it would be dishonest to pretend otherwise. When we ran the same experiment on home claimants, the surprise did not repeat. Even a claim handled well leaves a home customer more likely to look around than one who never claimed at all. The banked-trust story I have just told you is, on the evidence, a motor story. If I pooled the two products and gave you a blended number, both would be wrong and you would be misled. So home gets the next chapter to itself: why the product that pays out on the place a person lives obeys different rules, what a good claim is really worth there, and why the prize in home, though real and large, is a different prize altogether.

Chapter 5: The claim that hurts anyway

A crashed car goes away to be mended. A broken house stays exactly where it is, and you live in it.

That single difference, I have come to believe, explains almost everything the data is about to show you. A motor claim, for all its stress, happens somewhere else. The car is collected, or driven to an approved repairer, and the drama leaves with it. What remains is inconvenience: a courtesy car, a wait, a diary rearranged. A home claim happens to the place you sleep. The ceiling that came down is above your kitchen table. The dehumidifiers run all night outside your children's bedrooms, loud as small aircraft, for weeks. The claim does not go away to be fixed. You wake up inside it every morning until it is finished.

And it is rarely finished quickly. Take the peril that drives more home claims than almost any other: water escaping from where it should be, the burst pipe, the failed joint, the washing machine hose that let go while everyone was at work. A serious escape of water is not an event. It is a programme. Strip out, dry out, then repair, each stage with its own specialists and its own waiting list. The industry measures these claims in weeks and months. The customer measures them in family Christmases relocated and children doing homework in a bedroom that

smells of damp plaster. Storms and freezes pile the same misery up in bulk: in the worst recent years, weather pushed UK property claims to record levels, measured in billions of pounds. Behind every one of those pounds is a household living in the repair.

And a home claim arrives through many hands. This matters more than the industry realises. A motor claimant deals, mostly, with one machine: insurer, repairer, done. A home claimant may meet the call handler, the loss adjuster, the drying company, the builder, the builder's subcontractors, the contents specialist, sometimes a surveyor, occasionally a loss assessor they hired themselves in frustration. Every handover is a chance for the story to be retold, the promise to be re-made, the date to slip, the file to go quiet. In the last chapter I said a claim is a sequence of small promises. A home claim is a sequence of small promises made by strangers who have never met each other, performed in the customer's hallway.

So what does the data say? In the last chapter we ran an experiment on motor customers: hold the renewal price change constant, divide customers by claims history, watch who stays. Motor produced its surprise, the well-handled claim that beats never claiming at all. Naturally, we ran exactly the same experiment on home. And I have to give you the result straight.

Home is different, and it would be dishonest to pretend otherwise. A home claim is disruptive no matter how well it is handled. It happens in the place a person lives, it drags on for weeks, and it fills the house with strangers. Even a good home claim leaves customers somewhat more likely to shop around and switch than those who never claimed at all. Good handling does not beat never claiming in home. What it does, emphatically, is beat poor handling, by those twenty-six points. So the prize is real in both products, but it is not the same prize. In motor, a good claim can win you a customer outright. In home, it halves the damage of a moment your customer never wanted to have. Both are worth a great deal. They should never be sold as one.

Sit with the harder half of that first. Even when everything goes right, the home claimant emerges more likely to leave than the neighbour who never claimed. The banked-trust story of the last chapter does not repeat here. Why not? I do not think it is mysterious. The motor claimant's abiding memory is of a problem taken away and a promise kept. The home claimant's abiding memory, even after a good claim, is of the worst months in that house. The insurer may have performed well, and the customer may say so. But the policy itself has become a souvenir of a bad time. Renewal arrives and, fairly or not, some part of the customer wants a fresh start, the way people repaint a room after a burglary. You are not only competing with the market at that renewal. You are competing with the customer's wish to close a chapter.

Here the data offers a paradox, and it is one of the most instructive small findings in the whole survey. Home claimants are, if anything, slightly more satisfied with their claims than motor claimants. Eighty-three per cent against eighty-one. Read satisfaction alone and home looks like the healthier product. Yet it is the home claimant, including the satisfied one, who is more likely to walk. Satisfaction and loyalty, it turns out, are not the same thing, and home is where they come apart. Satisfaction is a verdict on you. Loyalty is a decision about the future, and it is taken in the shadow of the whole experience, including every part of it that was never in your control. A customer can honestly say you did well, mean it, score you five out of five, and still leave to put the whole episode behind them. Any measurement regime that reads satisfaction as a proxy for retention will therefore systematically misread home. Hold that thought; it returns when we talk about measurement properly.

There is one more wound home inflicts, and it is the quietest and perhaps the cruellest. It is the moment a family discovers, mid-claim, that the cover was not what they thought. The sum insured that will not rebuild the house. The contents limit that evaporates against a real inventory of a real family's possessions. The single item never listed. In the survey, when customers are asked why they went shopping for a new insurer, one of the answers they can give is exactly this, in their own words: the cover

was not what I expected when I claimed. And it is not a rare admission. Among home customers who went shopping, roughly one in four of those who had ever made a claim gave exactly that reason, however well or badly the claim itself was handled, against fewer than one in ten of those who had never claimed at all. A claim, of any kind, is where the cover gap comes to light. Think about where that discovery happens. Not at the quote, when it could have been fixed for a few pounds a month. At the claim, in the worst week, when nothing can be done. Technically, much underinsurance is the customer's own arithmetic. Emotionally, it lands as betrayal, and it lands on the denominator: you knew the rebuild costs, I did not, and you let me carry the gap. An industry serious about the home promise would treat underinsurance discovered at the claim as its own category of failure, one that no claims handler can repair, because it was committed at the point of sale, years earlier, by whoever let the promise be mis-sized.

So the home prize is different, but I want to insist, with numbers, that it is not smaller in the way that matters. The gap between a good home claim and a poor one is twenty-six points of a customer's future: roughly one in three good-claim customers walked over two years of renewals, against six in ten after a poor claim. That gap is nearly the size of motor's, and it survives the same price test. What differs is not the size of the lever. It is what the lever does. In motor, claims excellence is a growth weapon: it can leave customers more loyal than if nothing had ever happened. In home, claims excellence is damage control of the highest financial order: the difference between losing a third of your claimants and losing well over half. No other investment in a home book moves retention by anything close. It never turns the claim into a net positive, and a home insurer that promises its board otherwise is writing cheques the product cannot cash.

That reframes what a great home claim is even trying to do. If even a good claim hurts, the ambition cannot be delight. It is to make the hurt shorter, more predictable, and less lonely. Fewer hands, or at least one hand that visibly owns all the others. Silences filled before the customer

has to break them. Dates that hold. An end that arrives when promised, because in home, unlike motor, the customer is counting the days in their own hallway. Later in the book you will meet Mrs Davis, whose kitchen flood shows exactly what this looks like when a person, not a process, takes charge of it. I will let her story carry that argument when we get there.

Where does this leave us? With two products, two prizes, and one discipline: never sold as one, never measured as one. And with a question that the next chapter has to answer. Everything so far has compared the good claim with the poor one, the triumph with the disaster. But most claims are neither. Most claims are simply fine. The customer got paid, nothing went memorably wrong, nobody wrote a review. The industry looks at those claims and sees a job done. The data looks at them and sees something else entirely, and it is the most expensive thing in this book.

Chapter 6: Fine is not fine

Let me tell you about the most dangerous claim in your book. Nothing went wrong with it.

The customer rang. Someone answered, eventually. The forms were sent, filled in, processed. The assessor came, more or less when expected. The payment arrived, roughly when promised, for roughly the right amount. Nobody shouted. Nobody escalated. Nobody wrote a review, or a complaint, or a letter to the chief executive. Asked to score the experience, the customer thought for a second, said it was fine, and gave it a three out of five. Then they got on with their life.

Inside the insurer, that claim is a success. It closed on time. It cost what it should. It generated no complaint, no rework, no noise of any kind. On every dashboard the company watches, it is indistinguishable from a triumph. File it, and on to the next.

The tracker sees what happens next, and it is the most quietly expensive finding in this book. The customers who said their claim was fine go on to leave at very nearly the rate of the customers who said it was awful. In motor, roughly 56 in a hundred neutral claimants had gone within two years of renewals, against 58 in a hundred of those with a poor claim. In home, roughly half the neutral group left, against six in ten of the poor group. Look at those pairs of numbers for a moment. The claim that goes fine does not behave like a smaller version of the claim that goes well. It behaves like a slightly gentler version of the claim that goes badly. On

the loyalty ledger, fine sits next to failure, and nowhere near good.

I promised you this in the introduction, when I described the great, quiet middle a claims survey never hears from. Here is why the tracker hears them and almost nothing else does. Every other instrument the industry points at claims is tuned to intensity. A complaints process, by definition, hears only the angry. Reviews are written by the delighted and the furious, the people with a story burning to be told. Even the insurer's own post-claim survey is answered most eagerly at the extremes, and it is sent by the very company being judged, which colours who bothers to reply. The person whose claim was merely fine has no story to tell and no reason to tell it. They do not complain. They do not review. They do not answer. They just, at the next renewal or the one after, quietly go. The tracker catches them because it never asks about the claim first. It asks about their insurance in the round, and finds the claim along the way, with no axe to grind and no story required. The middle has always existed. It was simply inaudible.

And the middle is not small. Among motor claimants scoring their claim, for every customer who came away dissatisfied, there are nearly two who came away neutral. Sit those two facts side by side. The neutral group is nearly twice the size of the angry group, and it leaves at nearly the same rate. Multiply size by leak and the conclusion falls out on its own: most of the loyalty an insurer loses through claims does not walk out through the complaints department. It walks out through the middle, unrecorded, unexplained, and unmourned. And now the arithmetic has been run properly, the share can be put on it: of all the extra customers lost above the good-claim baseline, roughly six in ten are lost from the middle, in motor and in home alike. More loyalty leaves through indifference than through outright anger. You lose more customers to a shrug than to a scream. Anger is loud and rare. Indifference is silent and plentiful, and it is indifference that is draining the book.

Why should fine be so corrosive? Go back to the machinery this book has been assembling. A claim, we said, is the one moment the customer

finds out whether the promise was real. Chapter 4 showed what a good claim banks: proof, an answer the customer can hold against the price rise's question, what am I actually paying for? Now ask what a fine claim banks. Nothing. The promise was met, mechanically, and the trust question was never answered at all. Nobody took charge. Nobody made it feel easy. Nothing happened that the customer could not have imagined getting, slightly cheaper, from any name on a comparison site. The claim came and went, and the one opportunity the insurer will ever have to prove itself different was spent proving itself interchangeable. A poor claim answers the trust question with a no. A fine claim leaves it blank. And at renewal, blank and no turn out to cost nearly the same, because a customer with no proof has no reason. The fine claim is the dividend of chapter 3, forfeited. Not stolen. Just never collected.

There is a strategic consolation here, and it deserves naming, because this chapter should energise you rather than depress you. The angry tail is expensive to fix. Those claims usually went wrong somewhere real: a supplier failed, a decision was botched, a delay compounded. Recovering them means re-engineering whatever broke, and even then, some are unrecoverable. The middle is different. Nothing broke. Nothing needs apologising for. These claims were one phone call, one moment of visible ownership, one unprompted update away from being good, and the gap between fine and good is worth twenty-six points of loyalty in motor and sixteen in home. The middle is the largest pool of winnable loyalty in the book, and it is winnable with the cheapest tools in the business: attention, communication, and someone acting like they own the outcome. If I ran a claims operation and could fix only one cohort, I would not start with the disasters. I would start with the shrug.

Now to the other silence in the data, and I give you fair warning that this one has an ending I did not predict.

Turn the poor-claim number around. If nearly six in ten leave after a bad claim, then roughly four in ten stay. Stay, after the promise was tested and, in their own account, broken. Who does that? The comfortable

reading is loyalty: they weighed one bad episode against years of good history and forgave it. But there is a less comfortable reading, and the shape of this market makes it hard to dismiss on instinct alone. Perhaps they stayed the way people stay in a queue that has stopped moving. Not because they chose to, but because choosing costs effort, and effort is exactly what some customers, at some moments in their lives, do not have to spare. Renewal came, the letter sat on the side, and the path of least resistance renewed for them. That would not be loyalty. That would be inertia wearing loyalty's clothes, and on a dashboard the two are indistinguishable. Retention counts them identically. Only their reasons differ, and a retention figure cannot see a reason.

Some of these stayers, we can now say with numbers, are what I think of as silent detractors. They will never complain; we have already seen the middle does not complain, and the resigned complain even less. But they have not forgiven either. They are the customer who stays and tells the neighbour, over the fence, do not use them, you should have seen what happened to us. They sit inside the retention figure, propping it up, while working against the brand everywhere the brand cannot hear. A complaint count says these people do not exist. The recommendation data finds them, because you can stay with an insurer and still tell a researcher you would warn people away from it. And it finds them in extraordinary numbers. Among customers who stayed after a poor claim, seven in ten in motor, and two in three in home, score as detractors: they would actively warn others off the very insurer they still pay. Among good-claim stayers, the figure is one in twenty. The silent detractor is not a hypothesis. It is what most poor-claim stayers are.

So who, exactly, are the four in ten? Are they the busy and the comfortable, for whom an insurance premium is not worth an evening's admin? Or are they disproportionately the customers with the least room to manoeuvre, the older, the stretched, the ones least at home on comparison sites, the ones a regulator would call vulnerable? The survey holds the markers to test this, and we ran the test. The answer surprised me, and it is the happier one. The customers who stay after a

poor claim are not the ones who could not leave. If anything, they are more engaged with the market than the average customer who stays: more likely to have used a comparison site, more likely to have shopped around at that very renewal. They looked. They compared. And they chose to stay anyway. That is forgiveness, not inertia. On the markers the data can see, the queue theory fails: these are not people trapped by the effort of choosing, they are people who made a choice.

But hold the two findings of this section side by side, because together they say something subtler than either says alone. The poor-claim stayers chose to stay, and most of them would still warn a friend off. What they have forgiven, it seems, is the relationship, not the claim. They have given the insurer another year, and they are spending that year telling other people what happened. Forgiveness of that kind is not loyalty regained. It is a suspended sentence. And whether it also satisfies a regulator's idea of a market working, given who these customers are and what they carry, is a question with a different reader, later in this book. When we reach the regulator's chapter, we will put the same answer on the regulator's desk and see how it reads there.

For now, hold the two lessons of the quiet middle. First, the biggest leak in the book is not where the noise is; complaints are a map of anger, not a map of loss. Second, the cheapest loyalty available to any insurer is sitting in the claims that almost got there, waiting for someone to notice that fine was never fine.

Which completes the evidence, and leaves the sums. We now have the whole dividend laid out: what a good claim is worth against a poor one, in motor and in home. What a good claim banks against the price rise. What home really pays. What the indifferent middle forfeits. The pieces are all on the table, and each one has been priced in points of loyalty. The obvious next question is the one every finance director in every room has asked me, usually before I had finished the sentence. Fine. What does it come to in pounds? It is time to do the arithmetic, openly, with every assumption on the page.

Chapter 7: The arithmetic of the promise

In every room where I have presented the evidence you have just read, the same thing happens. Somewhere along the way, a finance director leans forward, waits for a gap in the sentence, and asks the only question their job allows them to ask. Fine. What does it come to in pounds?

This chapter is for that finance director. And the terms need setting before we start, because they are the terms that make the answer trustworthy. Everything that follows is arithmetic, not revelation. I will use round, stated assumptions for premiums, margins and claim rates, and I will show every one of them, because the moment a number's assumptions are hidden it stops being evidence and becomes marketing. The figures I derive are illustrations. Your book is not my illustration, and the real gift of this chapter is not any number in it. It is a method simple enough to redo on one sheet of paper, with your numbers, before your next board meeting.

The method has two lines. That is all.

The first line prices the premium. Take the loyalty gap from chapter 3, the difference in switching between a good claim and a poor one. Multiply it by the annual premium. The result is the premium a good claim keeps, per claimant, that a poor claim loses. Gap times premium

equals premium kept.

The second line prices the profit. When a customer walks, the business loses two things: the margin it would have earned on their renewal, and the acquisition cost it must now spend to replace them. So take the same gap and multiply it by the sum of those two: the acquisition cost avoided plus the margin on the premium retained. Gap times cost-of-replacement equals profit kept. Replacement cost belongs in the sum because retention and acquisition are the same pipe flowing in opposite directions; every customer a claim keeps is a customer marketing does not have to buy.

Now let us put numbers in, and watch what happens.

Start with motor. The gap, from chapter 3, is twenty-nine points, and remember what that number has already survived: it holds when you compare customers on the same price change, so no pricing artefact is hiding inside it. Assume an annual premium of £600, which is a deliberately unheroic figure for the current motor market. Twenty-nine per cent of £600 is roughly £171. That is the first output: each claimant whose claim is handled well rather than badly represents about £171 of premium kept at the next renewal, per year, at the same price.

For profit, assume it costs about £55 to acquire a motor customer, and take a net margin of five per cent, £30 on our £600 premium. Add them: £85 of value per retained customer. Twenty-nine per cent of that is roughly £24. So on these assumptions, the difference between handling one claim well and handling it badly is worth about £24 of bottom-line profit, and about £171 of revenue, per claimant, per renewal.

Home runs the same two lines with smaller inputs. The gap is twenty-six points. Assume a £300 premium: roughly £78 of premium kept per claimant. The same £55 acquisition cost plus £15 of margin gives £70 of replacement value, and twenty-six per cent of that is roughly £18 of profit per claimant.

I can feel the anticlimax from here. Twenty-four pounds? Eighteen? A

finance director does not rearrange a business for £24. But claims do not arrive one at a time, and this is where the arithmetic starts to compound into something a board cannot wave away. Take an insurer with a million motor policies. In any year, roughly six in a hundred of those customers make a claim. That is some sixty thousand claimants, each carrying that £171 of at-stake premium. Run the multiplication and the well-handled book keeps in the region of £10.6 million of premium a year, and about £1.5 million of profit, that the badly-handled book loses. Every year. Not from a new product, a new market, or a single pound of new customer acquisition. From the difference in how the same claims, from the same customers, at the same prices, were handled.

Scale it once more, to the whole market, and I do this only to show the size of the pool, not to promise anyone its contents. Across UK motor and home together, on the stated assumptions, the gap between good and poor claims handling is worth in the region of £500 million of premium a year, and something like £90 million of annual profit. Those are illustrative figures resting on every assumption I have named, and the appendix sets each one out with its base and its interval. But notice what they claim and what they do not. They do not say any firm will capture half a billion pounds. They say the prize pool that claims quality plays for, in this market, every year, is of that order. Industries restructure themselves for prize pools smaller than this one.

Now let me stress-test it in front of you, because a number you have not tried to break is not a number you should repeat. The most contestable assumption is margin. I used five per cent; the market has run leaner and fatter. Run the motor sums at three per cent and the market-level profit figure falls to roughly £39 million a year; at eight per cent it rises to about £55 million. Halve my assumed acquisition cost and the per-claimant profit drops by around a third; use the realistic acquisition costs of a price-comparison-driven market and it rises. Cut the claim rate, shrink the premium, apply whatever haircut your scepticism demands. What you will find, and what I have found in every room where someone tried, is that no plausible set of assumptions takes the answer anywhere near

zero. For the dividend to be worth nothing, the gap itself would have to be nothing, and the gap is the one input in this chapter that is not an assumption. It is measured, it is price-controlled, and it is the most robust finding in the book. The assumptions decide whether the prize is large or very large. They no longer decide whether it exists.

And everything so far is one renewal cycle, which understates the prize. The £171 is what the good claim keeps at the next renewal. But a kept customer is not kept once. They renew again, and their premium arrives again, every year they stay, while the replacement cost of the lost customer is spent only once but the lost premium is lost every year too. Value a retained claimant over a realistic tenure, even a modest one, and the per-claimant figures multiply. I am deliberately not printing a lifetime figure here, because it would pile assumption on assumption, and this chapter's authority rests on never doing that quietly. And there is a second, better reason for restraint, which is that the compounding has a clock on it. How hard the arithmetic compounds depends on how long a claim keeps working, and that is a question the data can now answer.

For years it was the most important open question in this book. Does the loyalty a good claim banks pay out at one renewal and fade, or does it carry for years? The survey can see this: claims sit in the data at one, two, three and four years' distance, and the gap can be cut by years-since-claim to draw what amounts to the half-life of a kept promise. So we drew it. The answer is decay, and it is steady. In the first year after the claim the gap is at its widest, some thirty-eight points in motor and thirty-three in home. By the second year it has roughly halved. In the third year it is still there, smaller but clean, around twenty points in motor and eighteen in home. By the fourth year it has faded into statistical noise, in both products. The kept promise, in other words, has a half-life, and it is measured in renewals, not decades.

The curve carries one more lesson, in why it falls. The gap does not close because the badly treated forgive; the poor-claim switch rate stays

high throughout. It closes because the well treated slowly drift back to the market. Year by year, the good-claim customer's switching creeps up towards, and eventually past, the rate of customers who never claimed at all. The banked trust is not withdrawn in one go. It evaporates, renewal by renewal, as the memory of being looked after fades and price resumes its ordinary pull. Nothing could say more clearly what the dividend actually is. What a good claim leaves behind is a memory, not a permanent upgrade to the customer, and memories need renewing.

So the right frame for this chapter's arithmetic is this. Persistence would have made the claim an asset that compounds by itself. Decay makes it an appointment. There is a window, two renewals, perhaps three, in which the trust a good claim banked is at full strength, and the insurer's job is to act inside that window: to renew the memory, to show up between claims, to convert the banked trust into the longer tenure that the one-cycle arithmetic only hints at. The per-claimant sums above are not diminished by this; year one, where they are calculated, is exactly where the gap is at its largest. What changes is the instruction that comes with them. The dividend is real, it is large, and it is perishable. A firm that handles a claim beautifully and then goes silent for three years has bought the window and wasted it. What the data cannot yet see is anything beyond the fourth year, because the survey's memory runs out there; the second edition of this book will extend the curve as the waves accumulate. But the question this chapter needed answered is answered. The claim is not an annuity. It is a moment, with a fuse.

Before I close the spreadsheet, one instruction, offered with affection to every finance director who has stayed with me this far. Do not take my numbers to your board. Take my method. It fits on a single slide, and the slide writes itself: our claimants this year, times our premium, times the gap, equals the premium we are choosing to keep or lose through claims handling. Your actuaries will argue about the inputs, and they should; the argument about inputs is precisely where the finance function starts taking claims seriously. What nobody in the room will be able to argue is the shape of it. The claim is a lever attached to real money, the lever is

measured, and someone in the organisation is currently deciding, mostly by accident, which way it is pulled.

Which leaves the question this whole part of the book has been building towards, and it should trouble you more the longer you sit with it. Here is a stream of value worth hundreds of millions of pounds a year across the market, per claimant sums any retailer would put on a dashboard by Monday, arithmetic simple enough for one slide. And yet almost no board in the industry has ever seen it. It appears in no statutory account, no monthly pack, no regulatory return. A dividend this large should be impossible to miss. So why has an entire industry, full of clever people who count everything, managed to miss it? The answer is not carelessness. The answer is that every instrument the industry uses to watch itself was built to see something else, and what a tool was not built to see, it cannot show. That is the part of the story you cannot see yet, and it is where we go next.

Chapter 8: The part of the story you cannot see yet

I left you, at the end of the last chapter, with a question that ought to trouble any leader in this industry. If a good claim is worth that much, why can almost no one see it? A dividend of between twenty-six and forty-plus percentage points of a customer's future, arithmetic that fits on one slide, sitting in plain view, and yet absent from the scorecards the industry studies every month. The answer is not that leaders are not looking. It is that the instruments they look through were never built to show it.

Every industry measures itself with the tools it has to hand, and over time it comes to believe that what the tools show is the whole picture. Insurance is no different. We have a set of well established measures for claims, and they are not bad measures. They are simply aimed at other questions. Ask them what they were built to answer, and they answer well. Ask them how a claim felt, and whether it built trust or destroyed it, and they fall silent. This chapter is about that silence, and about the value hiding inside it.

So let me name the problem plainly, once, because it is the pivot on which the rest of this book turns. The industry measures the claim it can count, and misses the trust that actually decides the customer's future. Almost everything the current view does well sits on one side of the trust equation.

Almost everything that matters sits on the other, unmeasured.

Take the measures in turn. We met the first already. Acceptance. Did you pay the claim. It is the industry's headline reassurance, and we have seen what is wrong with it. Almost everyone pays, at almost the same rate, so it cannot tell a great insurer from a poor one. It confirms the contract was honoured. It says nothing about how.

Then there are complaints. A firm watches its complaint numbers closely, and reasonably so. But complaints are a badly distorted mirror. Only a small and particular slice of unhappy customers ever complain. The rest do something quieter, and far more costly. They say nothing to you, and everything to the people they know. So a low complaint count is not the same as a contented customer base. It may just mean your dissatisfied claimants have stopped telling you and started telling their friends. Complaints catch the loudest unhappiness and miss the mass of it. And this is no longer merely my assertion. Lay the regulator's published complaints figures alongside what claimants themselves report, firm by firm, and the finding is this: the two correspond far more weakly than any board assumes. Knowing a firm's complaints rate tells you almost nothing about how satisfied its claimants actually are. A complaint count, it turns out, measures a customer's willingness to fight, not their likelihood to leave. The leavers, as we saw two chapters ago, mostly do not bother to fight.

Some firms go further, and survey their claimants directly after the event. That is a better instinct, and it reaches closer to the truth. But it carries a bias of its own, and a subtle one. The people who trouble to answer a survey about their claim are, overwhelmingly, the people who felt something strong. The delighted and the furious. The customer whose experience was simply acceptable, neither a triumph nor a disaster, rarely replies. We met these customers properly in chapter 6, and we priced what their silence costs: the quiet middle, nearly twice the size of the angry tail, leaving at close to the angry tail's rate. There is an old line that the opposite of love is not hate, but indifference. It holds

for customers too. The furious claimant is still engaged. Still telling you something. Still, sometimes, winnable. The indifferent one has stopped caring, and will drift away at the next renewal without a word of warning. A claims survey hears the love and the hate. It is deaf to the indifference. And chapter 6 has already shown you that the indifference is what costs you most, precisely because you never saw it coming. The point here is narrower and sharper: the instrument itself guarantees the blindness. Firms do not run their claims surveys badly. A claims survey run perfectly still cannot hear the people who decide the book's future.

And then there are the industry's published value measures. The figures on how often claims are made, how many are accepted, what they cost, how many complaints they draw. These do a useful job. They bring a welcome daylight to the market, and I would not be without them. But look closely at what they contain, and at what they do not. There is no measure of how the claim felt. No measure of whether the customer was treated as a person or a file. A firm can look entirely healthy on every published measure and still be losing the claimants who will decide its future. The measures were built to test fairness and price. They were never built to test trust.

Notice what all of these share. Acceptance, complaints, the value measures. Every one of them lives on the same side of the trust equation. They measure the parts a firm can count. Whether you paid. Whether you did what the rules require. The harder, countable terms, credibility and reliability. And every one of them is blind to the other side. To intimacy, whether the customer felt safe and cared for. And to self-orientation, whether they believed you were acting for them or for yourself. Those two terms are the ones that move loyalty. They are also the exact ground on which this industry is being accused of breaking faith. And nothing on the standard scorecard can see them at all.

Sit with that, because it is the heart of the matter. The whole public argument about whether insurance can be trusted is an argument about the denominator. Are you in this for me, or for yourself. And the

denominator is the one thing none of the industry's instruments measure. We are trying to defend our trustworthiness with a dashboard that cannot display the very dial in question.

You might reasonably say that at least a firm has its own claims satisfaction score. Its own direct measure of how its claimants felt. And that is the best instrument on the panel, by some distance. But even that one, the good one, has a hole in it. And the hole is larger and stranger than most leaders realise.

It comes from something peculiar to motor. When you have an accident that is not your fault, your claim is very often not handled by your own insurer at all. It is handled by the insurer of the driver who hit you. You are your own insurer's customer in name. But in the moment that matters most, the moment your car is wrecked and your week is upended, the experience is delivered by a company you never chose and will never renew with. Your own insurer, the one whose logo is on your policy, may barely touch it.

Now think about what that does to your satisfaction score. That claimant is living through one of the most formative experiences of their whole relationship with you. It will shape whether they trust you, and whether they stay. But as far as your claims system is concerned, they never made a claim. They are not in your numbers. Their experience, good or bad, happens entirely outside the one instrument you trusted to capture it.

And this is not a rare edge case. On the market data, around one in ten motor claimants had their claim handled by an insurer other than their own. A tenth of your claiming customers, living through a defining moment of the relationship, invisible to your scorecard. And when you can finally see them, as our data now can, they report a markedly worse experience than the claimants you handled yourself. So the hidden tenth is not merely unmeasured. It is, on average, among your least happy claimants, sitting in a blind spot precisely where the damage to loyalty is greatest.

And there is a subtler blindness still, and it explains why even a firm that

faithfully tracks what its claimants do next can miss the dividend entirely. Look at raw motor renewal data and good-claim customers appear to leave at almost exactly the rate of customers who never claimed at all. Three in ten, either way. A leader staring at that table would conclude, quite reasonably, that a good claim buys nothing. The table is misleading them. Good claimants, it turns out, disproportionately face price rises at their next renewal, and price rises drive switching. Only when you compare like with like, customers facing the same change in price, does the truth appear. The good claimants were the loyal ones all along. Their loyalty was masked by the prices they were being asked to swallow. The dividend was there the whole time. The instrument could not see it. This is the cruellest of all the blind spots, because it defeats the diligent. A firm could do everything this book has argued for so far, track its claimants to renewal, count who stayed, and still conclude from its own carefully built table that claims quality buys nothing. The raw number is not wrong, just answering a different question, and unless you know to hold price still before you read it, it will teach you the opposite of the truth.

I want to be careful how I put the next part, because it would be easy to hear all of this as an accusation, and it is not one. None of it is a failure of effort, or of care. It is a failure of instruments. The industry has been doing something hard, building trust at the worst moment of a customer's year, while looking through tools that cannot show the part that matters most. That is not a reason for shame. If anything, it is the opposite. It means there is a large amount of value here that no one has been able to see. And value you cannot see is value you have not yet lost. The moment you can see it, you can build on it.

So this is the doorway. We have established that the claim is the moment trust is decided. That a good one is worth more than almost anything else a firm can do. And now, that the tools the industry relies on are blind to the very terms that decide it, and blind, in part, even to some of its own claimants. The rest of this book is about closing that gap. About what a claim that builds trust actually looks like, moment by moment. About the people who deliver it, who become more important, not less, in the

years ahead. And about how to measure the thing that has been invisible, so that it can finally be led.

We begin where the trust is truly won or lost. Not in the scorecard, but in the journey itself. In the handful of moments, most of them small, where a frightened person decides whether they are in safe hands. That is where we turn next. And I want to start with a woman I will call Mrs Miggins, whose claim went perfectly, and who was not happy at all.

Chapter 9: The good claims journey

Let me tell you about Mrs Miggins.

She was seventy-five years old, and she made a motor claim. By every measure an insurer would normally use, it went perfectly. They took her car away and told her it would be back in a week. They gave her a replacement to drive in the meantime. And at the end of the week, exactly as promised, her car came back, repaired and beautifully cleaned. On time. On service. Everything correct.

So when we asked Mrs Miggins whether she had had a good claims experience, we expected a yes. She said no. She scored it poorly. And that puzzled us enough to go and find out why.

The replacement car, it turned out, was a little bigger than the one she was used to. She was frightened to drive it, so it sat on her drive, untouched. And the person who dropped it off had left it a foot over the edge of that drive, jutting out onto the pavement. For a week, Mrs Miggins felt she had to apologise to her neighbours for blocking the path. A kindness, the courtesy car, had become a source of worry and embarrassment.

Then we asked her what she actually used her car for. Two trips a week. One to the shops. One to the hairdresser. That was all. Mrs Miggins did not need a replacement car at all. What would have served her, what

would have made her feel looked after rather than burdened, was a taxi account, or simply the number of a local firm she could call.

Everything the process was designed to get right, it got right. And it still failed her. Because no one had asked the single question that mattered. Not what does the policy provide, but what does this person need.

I begin with Mrs Miggins because she dismantles a comfortable assumption, the idea that a good claim is simply a series of steps done correctly. Tick every box, hit every timescale, and satisfaction will follow. It does not. Trust is not built evenly across a claim, one increment for each completed task. It is built, or lost, in a handful of moments that carry far more weight than all the rest. Get those few right, and the customer will forgive a good deal of friction elsewhere. Get them wrong, and no amount of flawless administration will save you. Mrs Miggins had flawless administration. What she did not have was anyone attending to the moments that actually decide how a claim feels.

And here is the pattern that matters. Those moments, the ones that move loyalty, are almost never the ones you would find on a process map. They are not about speed, or paperwork, or the mechanics of the repair. They are about how the person was made to feel. Whether they felt understood. Whether they felt safe. Whether, at the worst moment of their year, they had the sense that someone was on their side. This is the term of the trust equation I have said least about so far, and it is the one this chapter belongs to. Intimacy. Not intimacy in any grand sense. Simply the felt experience of being cared for by someone who sees you as a person rather than a case.

Remember who is on the other end of a claim. Not a policy number. A person, usually at a low ebb. They have pranged their car, or watched their home flood, or come back to find their things gone. They are shaken, sometimes frightened, often embarrassed, rarely at their best. They are not in the mood to admire your efficiency. They want to know they are in safe hands. Every moment of the claim is really them asking that one question, over and over. Am I safe here. Are you going to look after me.

And they read the answer not from your service levels, but from how they are treated.

When you look across the journey for the places that answer that question loudest, a few stand out.

The first is the very first contact. The moment the person makes the call, or opens the app, and tells you what has happened. They are frightened, and they do not yet know whether this is going to be a fight. Whether they are met with calm and reassurance, or with a script and a reference number, sets the emotional temperature for everything that follows. A great first contact says, in effect, you have done the right thing, we have got you now. Very little else in the claim matters as much as that first breath of relief, or its absence.

The second is what happens in the silences. The stretches between updates, when the customer is left waiting and wondering. This is where more trust quietly leaks away than almost anywhere else. Not because anything has gone wrong, but because no one has said anything. Being kept in mind, hearing from you before you have to be chased, is one of the plainest signals of care there is. Its absence is one of the loudest signals of indifference.

The third is the moment Mrs Miggins never got. The moment someone looks past the policy and asks what this particular person actually needs. It costs almost nothing. A question, and the willingness to act on the answer. But it is the difference between a customer who feels processed and one who feels seen. It is, in many ways, the whole of intimacy compressed into a single decision.

Now, I have just made three confident claims about where a claim is won and lost, and for most of my career, statements like these could only ever be offered as practised instinct. Believe me, I have sat through enough conference talks built on nothing firmer. So let me hold my own argument to the standard this book has set. The survey behind this book does not only ask claimants for an overall verdict. It walks them back through the claim itself, a battery of specific statements, eleven

in motor and nine in home, each rated: how easy it was to make the claim and to get in touch at the outset. Whether they were kept fully informed. Whether they knew what to expect next. Whether they knew how to reach their insurer throughout. Whether the claim was handled sympathetically. Whether the staff were friendly and helpful. Whether it was resolved quickly. And then each product's own tail: in home, whether the outcome was right; in motor, the physical stages, the repairer, the quality of the repair, the courtesy car. Mrs Miggins's stage is in there, measured across the market. Which means the three moments I have just described are not beyond evidence. Each one maps onto statements in that battery: the first contact onto ease of getting in touch, the silences onto kept informed and knew what to expect, being seen onto sympathy and human warmth. And the statements can be ranked, coldly and fairly, by a single test: how far does each one separate the claimants who stayed from the claimants who left?

That test has now been run, and every statement on the list passes it: each one separates the stayers from the leavers by a margin the statistics stand behind, in both products. Here is the full ranking. At the top of motor's list sits the plainest promise in the battery. Did the customer know what to expect through the process? Claimants who agreed were twenty-seven points more likely to stay with their insurer than those who did not. Being kept fully informed runs second, at twenty-four points. The friendliness of the staff runs third. In home the order reshuffles: the ease of making the claim in the first place leads the list, with speed of resolution and knowing what to expect just behind. Read that against my three moments, and the data is kind to them, though it insists on its own emphasis. The silences, knowing what is coming and being kept informed, carry the top of the motor table. The first contact carries the top of home's. The moment of being seen, the sympathy and the warmth, sits in the body of both lists rather than at the head of either, but it separates stayers from leavers decisively everywhere it is measured, and, as we are about to see, it carries as much of a customer's future in the round as anything the process can deliver. The three moments survive the test. Which of them

matters most depends on the product, and none of them may be skipped.

And this is where the promissory note from chapter 3 gets paid. I said there that reliability was, provisionally, the term that most dependably converts a claim into a kept customer, and that I would test the belief rather than assert it. The statement-level analysis maps each stage onto the trust equation, kept informed and knew what to expect and resolved quickly to reliability, sympathy and friendliness to intimacy, and lets the switching data crown the winner. It crowned no one. The reliability statements and the intimacy statements separate the stayers from the leavers by almost exactly the same amount, in motor and in home alike. In home, the single strongest statement is not even a reliability item; it is the ease of that first contact. I could not be happier to have been wrong. Because look at what the dead heat means. It means the human half of a claim, the warmth, the sympathy, the sense of being in kind hands, is not the garnish on an efficient process. It is, measurably, half the dividend. A claims operation that keeps every timescale and engineers out the humanity forfeits as much of a customer's future as one that is warm and shambolic. The industry has spent years debating which to invest in, the machinery or the people, as though the data would one day settle the argument by picking a side. The data has now spoken, and its answer is that the argument itself was the mistake. Mrs Miggins was never just an anecdote. She is what the whole market's numbers say: the process can be perfect, and if the person is not seen, the claim is not won.

And running through all of them is the handler. The person, or the handful of people, the customer actually deals with. Because every one of these moments is delivered by a human being, and felt by another. To the customer, in that moment, the handler is not an employee of the company. The handler is the company. The entire brand, everything the marketing ever promised, is standing there in the voice of one person on the phone. That is an extraordinary thing to place on someone's shoulders, and we will come back to what it asks of leadership.

If Mrs Miggins shows what happens when no one thinks to ask, let me

give you a companion who shows something almost worse. Mrs Davis spelled it out, and it still did not land.

Mrs Davis had an escape of water at home. Anyone who has handled home claims knows what that means. It is not one visit but many, and many different trades. People to dry the place out. Plasterers. Someone for the floors, someone for the ceilings, painters, decorators. A parade of separate crafts, each arriving in turn across the weeks it takes to put a home back together.

When the first loss adjuster arrived, he rang the doorbell, and Mrs Davis's dogs went wild. So she made one small, clear request. Please do not ring the doorbell. Please make sure nobody rings it. Come round to the back and knock on the back door instead, because the bell sets the dogs off, and that distresses me. It could not have been simpler, or more reasonable.

Eight different trades visited that house over the course of the claim. Every single one of them rang the front doorbell.

Nothing about this was hard. It was the easiest request in the world to honour. It failed for one reason only. No one carried it. The request sat with the first person who heard it and travelled no further, because the claim was run as a list of tasks to be completed, not as a person to be looked after. Did the claim get settled? Yes. Was the water damage put right? Yes. Did the dogs bark, eight separate times? Yes. Was Mrs Davis upset, eight separate times? Yes.

That is the thing about intimacy in a complex claim. It is not enough for one kind handler to get it right. In a home claim especially, the customer is passed from hand to hand, and every hand is, to her, the company. The care has to survive the handoffs. It has to be noted, passed on, and honoured by the eighth visitor as faithfully as the first. Each of Mrs Davis's eight trades did its job. Not one of them did hers.

Which brings me to the sentence I most want you to take from this chapter. Claims are not processes. They are people. The process matters, of course it does, and later I will argue that getting it right is what frees

people to do the part only they can do. But the process is the floor, not the ceiling. Mrs Miggins proves it. Her process was immaculate and her experience was poor, because the moments that decide a claim are human moments, and no process, however well designed, can feel for you. It can only clear the way for a person to.

So what does this ask of you, if you lead a claims operation. Less than you might fear, and more than a process review will ever give you. It asks that you design for the moments that matter, and not only for the steps you can measure. That you find the few points in the journey where trust is actually won, the first contact, the silences, the moment of understanding, and pour your attention there. And it asks that you give your people the room, and the permission, to be human at exactly those points. To ask the extra question. To offer the taxi account instead of the car nobody wanted. To behave, at the moment of truth, like one person looking after another.

None of this is soft. These moments are where the dividend of the earlier chapters is actually won, one claimant at a time. But you will have noticed a difficulty, and it is the one chapter 8 left us with. Moments like these are almost invisible to a standard scorecard. You cannot manage what you cannot see, and the very things that matter most here are the hardest to measure. That is the problem we turn to next. Because getting these moments right, consistently, across thousands of claims, is not a matter of luck, or of individual kindness. It is a matter of what an organisation chooses to value, and of leaders who can see the moments clearly enough to build around them. First, then, we have to make the invisible visible. We have to learn to measure the thing that matters.

Chapter 10: Measuring what matters

The last chapter closed on a promise and a problem. The moments that decide a claim are human moments, the ones Mrs Miggins and Mrs Davis never got. And human moments are precisely the hardest things to measure. So the promise, to make the invisible visible, sounds close to impossible. This chapter is about how it is done.

The first obstacle is a belief, held across much of the industry, that the things which matter most in a claim cannot really be measured. That feelings are soft. That trust and intimacy are too vague to count. And so we fall back on measuring the hard things, the timescales and the costs, because at least those hold still. I want to take that belief apart, because it is the single thing standing between the industry and the dividend we have been describing. Trust can be measured. Not perfectly, but well enough to manage. You have to measure it the right way, from the right place.

And the right place is the only place the felt experience actually exists. The customer's side. You cannot see intimacy from inside the process, because intimacy is not a step in the process. It is how the whole thing landed on a frightened person at their kitchen table. The only way to know whether Mrs Davis felt looked after is to ask Mrs Davis. So the first principle of measuring what matters is almost embarrassingly simple.

Ask the people. Ask them properly, and ask them at scale.

But asking is not enough on its own, because who you hear from decides what you learn. We saw the trap earlier. Ask only in the aftermath of a claim, in a survey badged as a claims survey, and you hear from the delighted and the furious, and almost no one in between. A measure built on that is worse than no measure, because it tells you a confident story that is not true. To measure what matters, you have to hear the whole spread. The quiet middle. The hidden tenth whose claim was handled by someone else. The people who would never trouble to fill in a form about their claim, but who will tell you the truth if you reach them another way. A measure you can trust is one that hears everyone, not only the ones with a story to tell.

The third principle is the one that changes everything, and it is the one a single insurer can never satisfy alone. A number on its own means almost nothing. Suppose I tell you that eight in ten of your claimants were satisfied. Is that good? You have no idea, and neither do I, until I can tell you what everyone else scored. Eight in ten is excellent if the market sits at seven. It is alarming if the market sits at nine. A claims score has no meaning in isolation. It only has meaning against a line. So the foundation of the whole method is a market line. A view of where the entire market sits, built the same way for every firm, against which any single firm can at last be read. And here is the crucial part. No insurer can draw that line for itself, however good its own data, because no insurer can see anyone's claimants but its own. The line can only be drawn from outside, from a vantage point that takes in the whole market at once. That is what makes it valuable. And it is what the industry has never had until recently.

There is a subtlety in drawing that line, and it is worth a paragraph, because it decides whether the line tells the truth. Customers choose brands. They buy from the names they know, and it is the brand they judge when the claim comes. But behind the brands, the market is more concentrated than customers realise: one firm, one claims operation,

one machine, can sit behind several of the names on the comparison site, handling all of their claims. So whose experience are you measuring? The answer is that the line has to be readable both ways. Read it by brand, and you see what customers actually chose and felt. Read it by the firm behind the brands, and you see what one claims machine is really delivering across everything it touches, which is the view that matters for running it. A measure that can move between those two views can catch what neither view alone would show: the single machine producing different experiences under different names, or one weak brand dragging on an otherwise strong operation. The regulator's published data, incidentally, can only ever be read one way, by the authorised firm. The customer's account of the claim is the only vantage point that can see both the brand that was promised and the machine that delivered.

The fourth principle keeps the whole thing honest, and keeps it commercial. A measure of experience that floats free of consequence is just a satisfaction score, and boards are right to be wary of those. The measure has to be tied to what the experience is worth. It has to connect the felt moment forward to the thing the business cares about, which is whether the customer stays. When you can trace a claim experience through to loyalty, and loyalty through to growth, satisfaction stops being soft. It becomes the leading indicator of retention, which is to say, of the future of the book. And the tie-through has to clear one final bar, because there will always be a voice at the table, usually a well-paid one, saying that loyalty is really just pricing wearing a costume. So the connection between experience and staying must be shown to survive a price control, compared like for like among customers who saw the same change in price. As we saw in chapter 3, it does, and barely narrows. That is what closes the argument in the boardroom. A sceptical board can dismiss a satisfaction score. It cannot dismiss a satisfaction score that predicts retention after price has been held still, because at that point the measure is no longer describing feelings. It is describing money that has not arrived yet.

There is a fifth principle, and it is the one that turns a measure from a verdict into a plan. A score on its own tells you where you stand. It does not tell you why you are there, and without the why you cannot know what to change. This is the quiet failing of most satisfaction measurement. A customer is asked to rate their claim from one to five, they choose a two, and the number is dutifully recorded. But no one knows what the two means. Was it the long silence in the middle? The replacement car that did not fit? The eighth person to ring the doorbell? The score captures the symptom and loses the cause. A measure worth having has to carry both. The what, and the why. It has to show not only that a claim fell short, but where, and in which of the moments that decide how a claim feels. Only then does a number become something a leader can actually act on.

Step back from these principles and notice what they add up to, because it is the heart of this chapter. Every insurer in the country will tell you it is good at claims. It is on the websites and in the advertising. We are here for you. We will look after you. But the customer has learned to discount all of it, because it is the company praising itself, and self-praise is the cheapest thing there is. Saying you are good at claims is an act of self-orientation, and the customer hears it as one. Proving it is something else entirely. This is the term of the trust equation this chapter belongs to. Credibility. Credibility is the difference between claiming and showing. And a measure of this kind, independent, market-wide, drawn from the claimants themselves, is a credibility instrument before it is anything else. It lets a firm stop saying trust us, and start saying here is where we stand, measured exactly as everyone else is measured. That is a quieter sentence. It is also a far more powerful one.

A book like this can carry the market line. It can show you where the market sits, how wide the spread runs, what a good claim is worth. All of that belongs in these pages, freely, because the whole industry is better for knowing it. What a book cannot carry is the part that is yours alone. Your own position on that line, and, more to the point, your own why. The particular moments where your claims win trust or lose it. That is

specific to you, drawn from your own claimants set against everyone else's, and no general book could ever print it. I raise it only to make one point. The what and the why are knowable now, in a way they were not a few years ago. What any firm chooses to do with that is its own affair.

And if you let the idea run a little further, something larger comes into view. If claims quality could be measured the same way across the whole market, against one shared line, then the industry would have, for the first time, a common language for the thing it has always struggled to talk about. Not each firm marking its own homework, but a single yardstick everyone could point to. I believe that is where this ought to go, and I will come back to it near the end of the book, because it turns out to be as much a question of leadership as of measurement.

Because a measure, however good, changes nothing on its own. You can know precisely where you stand and do nothing about it. A number on a page has never improved a single claim. What turns measurement into better claims is what someone decides to do with it. And that is leadership. Not leadership as a slogan, but the practical, daily work of building an operation that acts on what it can finally see. It is, I will argue, the hardest and most important part of this whole story. And it has two halves that most organisations keep apart. Process, and people. Learning to hold them together is the work of the next chapter.

Chapter 11: Where process meets people

The last chapter ended on a hard truth. A measure, however good, changes nothing on its own. You can know exactly where you stand, and exactly why, and still do nothing about it. What turns knowledge into better claims is leadership. And claims leadership, it turns out, is a particular kind of work, with two halves that most organisations keep in separate rooms.

Every leader I have watched succeed in claims has been good at two very different things, and it is the combination that is rare. The first is process. The second is people. Most leaders lean naturally towards one. The great ones refuse to choose, and hold both. This chapter is about why they must.

Start with process, because it is the half the industry understands best. Process leadership is the skilled work of seeing how a claim actually flows. Knowing where the steps are, where they collide, where the friction and the delay and the cost pile up. And then driving efficiency and effectiveness through the whole of it. This is real and demanding work, and it is not to be dismissed. Done well, it is what keeps the promises a customer can feel. The assessor who comes when promised. The repair that holds. The update that arrives on time. In the language of the trust equation, process leadership is how you deliver credibility

and reliability at scale. The competence, and the kept word.

But process has a failure mode, and we have already met it. Left to itself, process optimises the transaction and forgets the person. It becomes very good at the thing it can measure, and blind to the thing it cannot. Mrs Miggins is what process leadership looks like without the other half. A machine tuned to return a car in a week, performing flawlessly, and never once asking whether the car was what she needed. Efficient, and hollow. Process alone can hit every target and still leave the customer cold.

The second half is people, and it is the one the industry finds harder to talk about, because it is harder to count. People leadership is about values, and behaviours, and how the people who deliver your claims actually feel. It matters for a reason that is easy to miss. A handler who feels trusted, valued, and clear about what the company stands for behaves differently towards a frightened customer than one who feels squeezed and unseen. And here is the part I have come to believe most firmly, from years of listening to customers. They can tell the difference. They feel it down the phone. They could not always name it, but they know whether they are being dealt with by someone who has room to care, or someone processing them under pressure. A good claim and a bad claim often trace back not to procedure, but to values. To how the organisation treats the people who treat the customer.

Yet people alone is not the answer either, and it would be dishonest to pretend otherwise. Warmth without process is a lottery. Wonderful with one handler, absent with the next, because nothing holds it in place. It cannot be relied upon, and it cannot be scaled, and over time it becomes expensive, in money and in the good people who burn out trying to care inside a system that does not help them. Mrs Davis is what people leadership looks like without the other half. The care was real. It existed in the first person who heard her ask about the doorbell. And it died at the very first handoff, because there was no process to carry it to the second visitor, or the eighth. Warmth with nothing to hold it does not survive contact with scale.

Put the two halves side by side, and the trust equation sorts itself almost perfectly between them. The countable terms, credibility and reliability, are won by process leadership. The human terms, intimacy and the sense that someone is on your side, are won by people leadership. Neither leader, working alone, can move all four. And all four is what trust requires. Which is why the equation is, in the end, a description of a leadership job. Tend the top line with your process. Tend it also with your people. And keep your eye, above all, on the number underneath. The journey data of chapter 9 now shows this split in numbers, and the numbers refuse to choose. The stages process owns, the repair, the speed, the mechanics, and the stages people own, the keeping informed, the sympathy, the warmth, each carry almost exactly the same share of a customer's future. The data has settled the argument between the two camps by dismissing it.

Because the denominator, self-orientation, is set almost entirely by people leadership. Whether a customer believes you are acting for them or for yourself is not decided in the moment of the claim. It is decided long before, in the culture that shapes how your people see the customer. A claims operation run purely to contain cost teaches its people, quietly, that the customer comes second, and customers feel that orientation even when no one ever says it aloud. An operation built around genuine care teaches the opposite. The most important term in the whole equation is not a process output. It is a cultural one. And culture is the truest test of a leader.

So the answer is not to choose a side, and yet choosing is the constant temptation. There is an efficiency camp that worships process and treats empathy as a soft indulgence. There is an experience camp that worships empathy and treats process as the enemy of care. Each is half right, which is exactly what makes each one dangerous. The truth is that trust needs both, and that the two live in real tension. Empathy costs the very time that efficiency wants to save. Efficiency can flatten the very moment empathy needs to breathe. The leader's real work is not to resolve that tension by picking a winner. It is to hold it. To build process that frees

people to care, rather than process that stops them. And to grow a culture that makes people want to, rather than one that grinds it out of them. That fusion, process and people held together under tension, is the whole of great claims leadership. It is also hard, which is why it is rare, and why those who manage it deserve a great deal more credit than the industry tends to give them.

If you lead claims, you already know this in your bones, because you live the tension every day. You are asked to be fast and kind at once, careful with cost and generous with care, at the worst moment of your customer's year, thousands of times over. That is one of the hardest leadership tasks in any industry. This book is, in large part, an argument that it is also one of the most valuable, and that the person who masters it is building something no competitor can easily take.

And now a new force is arriving that will test the fusion as nothing has before. Artificial intelligence is coming for the process half of claims, and it is very good at it. Faster, more consistent, more tireless than any human operation. To many that will sound like the argument settled at last in favour of efficiency. I want to persuade you of the opposite. But I cannot do that honestly without first showing you the world the machine is arriving into, because it is not arriving into a calm industry with money to spare. It is arriving into an industry under genuine strain: claims costs rising faster than premiums, fraud that has become organised and industrial, results that leave little room for generosity. That strain is real, it deserves a fair hearing, and it will shape what the machine is asked to be. A squeezed industry reaches for a cost tool. So the next two chapters travel together. First the pressure, told fairly from the industry's side of the ledger. Then the machine, and the choice it forces. Because whether AI ends up industrialising care, or industrialising its absence, will be decided not by the technology, but by leaders holding exactly the fusion this chapter has described.

Chapter 12: The pressure on the promise

I have spent eleven chapters telling you what customers deserve at the moment of a claim. Fairness demands I now spend one telling you what the people delivering that moment are up against. Because if you have read this far as a claims leader, somewhere around chapter 7 a voice in your head will have started saying: he is pricing the dividend as if I had money to chase it with. So let me show you the other side of the ledger, with no softening for the sake of the argument. This book promised at the outset to be a champion of this industry rather than a critic. Champions do not earn that name by ignoring the weight on the bar.

Start with the plainest fact of the current market, the one that frames everything else. In motor, the forecasters expect the industry to spend more paying and running claims than it collects in premiums. The projected combined ratio, claims plus costs as a share of premium, sits comfortably above one hundred per cent: on recent forecasts, somewhere between £1.07 and £1.11 going out for every £1.00 coming in. Dwell on what that means. The core activity of insuring cars, taken as a whole, is currently done at an underwriting loss. Firms survive on investment income, on ancillary revenue, and on hope, and a market that runs at a loss for long enough produces exits, failures and consolidation, all of which are already visible. This is not a market sitting on a cushion,

idly declining to invest in claims experience. It is a market where the claims line is drowning the account.

And the claims line is not drowning because insurers have become careless. It is drowning because the underlying cost of putting things right has been rising structurally, from several directions at once, and most of them are nobody's fault.

Consider what has happened to the car itself. A windscreen used to be a sheet of glass; now it is a mounting for cameras and sensors, and replacing it means recalibrating the driver-assistance systems that look through it, turning what was once a modest bill into a specialist's job. The technology that prevents accidents makes each surviving accident dearer, a strange bargain that the public has never been asked to notice. Electric vehicles push the same direction harder: more complex to repair, needing specialist handling, and carrying a battery that, if damaged, can turn a repairable car into a total loss on its own. Around the vehicle, the repair trade itself has been squeezed by parts and paint inflation and a shortage of skilled hands in the bodyshops, which raises prices and stretches the very timescales that chapter 9 showed customers judging us by. When the claim involves a third party, the costs compound: the replacement car supplied on credit hire while a liability argument grinds on, billed at rates that would make a customer blink. And in the gravest cases, the ones involving serious injury, settlement costs are shaped by discount-rate rules and injury-reform outcomes that can move an insurer's liabilities by enormous sums at the stroke of a pen. Stack it together and the volume is remarkable: in a single recent quarter, UK motor insurers paid out on the order of £3.2 billion across some six hundred and thirty thousand claims, record territory.

Home tells the same story with different weather. We saw in chapter 5 what an escape of water does to a family; multiply it by every burst pipe in a cold snap and add the named storms, and recent years have pushed UK property claim payouts to record levels, measured in billions per year. Climate does not negotiate, reinsurance for it does not get cheaper, and

every flooded winter arrives on the claims line of an account that was priced the previous spring.

Then there is the pressure that comes not from physics or weather but from deceit. Fraud against UK insurers runs, on the industry's own detection figures, at something over a billion pounds a year, and the direction of travel is worse. Two developments make it nastier than the raw number suggests. The first is organisation: much of it is no longer the opportunist rounding up a genuine loss, but networks, staged accidents, ghost brokers selling fake policies to real people. The second is newer and should trouble everyone: the same artificial intelligence arriving in claims departments is arriving in the fraudster's toolkit, and it has collapsed the cost of producing a convincing lie. A fabricated invoice, a doctored photograph of damage, a plausible document trail: what once took skill now takes a prompt. One major insurer alone reported detecting a third more fraud in a recent half-year than the year before, thousands of cases, tens of millions of pounds. And remember that all detected fraud is paid for by someone, and the someone is every honest customer's premium.

Now put yourself, properly, inside the building that has to hold all of this. You run a claims operation. Your severity inflation outruns your pricing. Your combined ratio starts with a one-one. Your board wants the claims line contained, because it is the largest line there is. Your regulator, quite rightly, requires fair value and good outcomes and will name you if you fail. And a meaningful slice of the claims coming through your door this morning are lies, some of them industrially produced, and you cannot always tell which. Every pound you spend, someone upstairs would like back. Every hour your handlers spend caring, the efficiency programme would like back too. This is the simultaneous equation the people of this industry are solving every day, and it is why I have no patience with commentary that treats claims departments as villains. The people I know in those buildings are doing something authentically difficult, under pressure from every direction at once, including from books like this one.

But I owe you the tension at the heart of it, because it is the hinge of this chapter and it cannot be resolved by goodwill alone. Every fraud control that protects the loss ratio adds friction for the honest claimant. Every one. The extra verification step, the second set of photographs, the request for receipts a flooded family cannot produce, the pause while the counter-fraud model scores the case: each is rational, each is defensible, and each is experienced by the honest customer as suspicion. And here is the arithmetic that makes it cruel. The overwhelming majority of claimants are honest. Controls calibrated to catch the dishonest few are felt by the honest many, at the precise moment this book has spent eleven chapters proving to be the hinge of the entire relationship. Recall the denominator. The customer's silent question at the claim is, were you on my side or your own? A fraud control, however necessary, whispers the wrong answer. The claims leader is thus asked to hold two truths about every single claimant simultaneously: this person is probably a frightened customer at the worst moment of their year, and this claim might be a fabrication. Treat the fraudster as a customer and you leak money. Treat the customer as a fraudster and you leak everything this book has priced. There is no clean escape from that tension, only the daily craft of holding it, which is one more reason chapter 11 called claims leadership one of the hardest jobs in any industry.

I should also be straight about what this pressure does to the dividend argument. It does not weaken the arithmetic; the gap survives every assumption we threw at it in chapter 7. What the pressure does is explain, better than any theory of neglect, why the dividend goes unharvested. Generosity of attention feels unaffordable in a account running at a loss. Investment in experience competes, line by line, against costs that are rising on their own. A business drowning in severity inflation does not lift its eyes to retention economics, for the same reason a man bailing a boat does not stop to varnish it. The tragedy, and by now you can see it coming, is that the varnish in this metaphor is holding the boat together. The premium kept by a good claim is exactly the revenue a loss-making account most needs, and it is the line of defence that costs least, because,

as chapter 9 showed, the widest gaps between good and bad sit in the cheap human stages, the keeping informed, the being kind, not in the expensive metal ones.

Which brings us to the door this chapter was always walking towards. Put yourself back in that building one last time. Costs rising structurally. Fraud industrialising. A regulator watching. A board demanding containment. And onto your desk arrives a technology that promises, credibly, to handle claims faster, cheaper, more consistently, around the clock, while scoring every case for fraud as it goes. Of course the industry will reach for it. It would be irresponsible not to. A squeezed industry meets a cost-cutting machine, and the attraction is not a mystery; it is arithmetic. The only question that matters, and it is the question the next chapter exists to answer, is which half of the claim the machine will be pointed at. Everything this book has found says the claim is two things at once: a process that should be as fast and clean as engineering can make it, and a human moment that decides the customer's future. A machine pointed at the first is the best news claims has had in decades. A machine pointed carelessly at the second will industrialise indifference, at scale, with perfect consistency, and chapter 6 has already shown you what indifference costs. The pressure is real. The machine is coming. What happens next is a choice, and it is to that choice we now turn.

Chapter 13: The machine and the promise

The last chapter closed with a promise to change your mind. Artificial intelligence is arriving in claims, and the common assumption is that it settles the old argument between efficiency and care once and for all, in favour of efficiency. I believe it does the opposite. To see why, we have to be clear-eyed about what the machine is very good at, and clear about what it can never be.

Let me concede its power first, because there is no sense pretending otherwise. AI is going to be extraordinary at the process half of claims. It can respond in the first second rather than the first hour. It can keep a customer informed without ever forgetting, chase every open task without tiring, and hold a thousand timescales in mind at once. The update that always arrives. The promise always kept. The status always known. In the language of this book, AI can deliver credibility and reliability at a scale and a consistency no human operation could ever match. That is a gift, not a threat, and any leader who refuses it will be beaten by one who does not.

But now hold the trust equation up against the machine, and something important appears. AI is superb at the two terms it can automate, and powerless at the two that decide trust. It cannot deliver intimacy, because intimacy is one human being feeling understood by another,

and a frightened person at the worst moment of their year knows the difference between being handled by someone who cares and being processed by something that cannot. And it does something more dangerous still, down at the denominator. A machine visibly optimising a claim looks exactly like an insurer serving itself. When a customer senses that an algorithm is deciding their outcome, the oldest suspicion in insurance is right there waiting. This thing has been built to pay me as little as possible. Left to run the whole claim, AI can raise efficiency while deepening the very wound the industry most needs to heal.

So the obvious path is a trap. Automate everything, strip out the cost, let the machine run the claim from end to end. It wins the cheap, countable terms and loses the expensive, decisive ones. It is Mrs Miggins all over again, only now at industrial scale. A flawless operation that returns the car on time and never once asks what the person needs. The efficiency camp will call it progress. The customer will call it the moment they stopped trusting you.

The answer is not to keep people away from the machine, nor the machine away from people. It is to use each for what only it can do. Let the machine take the process. Let the human take the trust. Give the automation every routine task it can carry, the tracking and the chasing and the updating, precisely so that your people are freed from the drudgery that was swallowing their attention. And then spend the hours you have freed on the handful of moments that actually decide the claim. The first frightened call. The silence that needs breaking. The question no algorithm thinks to ask. Automate the process, and reinvest the humanity. That is the whole strategy, and it is the fusion of the last chapter made suddenly affordable.

Which leads to the conclusion I most want you to carry out of this book, because it runs against almost everything you will hear about machines and work. In claims, AI does not make the human handler less valuable. It makes them more valuable. As the routine is absorbed, the scarce and precious thing becomes the human moment, and the person who can

deliver it. The handler stops being a processor of tasks, which is what the machine is for, and becomes the carrier of trust, which is what the machine can never be. Remember that to the customer, the handler is the company. In a world where the process is automated and near identical from one insurer to the next, that human presence is no longer a cost to be minimised. It is very nearly the only thing left that one insurer can offer and another cannot. It becomes more of the promise, not less.

And notice that none of this is decided by the technology. The same tools can be used to remove people in the name of savings, or to free people in the name of trust. The machine does not choose. The leader does. Which future an insurer walks into is not a question of what the software can do. It is a question of what the leadership values, which is exactly where the last chapter left us. The denominator, once again, is set in the culture, not in the code.

But a choice like this can only be made well if you can see what you are trading. Automate a step, and did trust rise, or quietly fall? Move a moment from a person to a machine, and did the customer feel freed, or abandoned? You cannot steer by cost alone, because cost will always argue for the machine. You need to see the human terms, in the round, and against the market, or you will optimise your way into Mrs Miggins without ever noticing. This is why the measure matters more in the age of AI, not less. It is the only instrument that can tell a leader whether the machine is building trust or spending it.

And here, near the end, I want to lift my eyes past the single firm, because there is a larger prize in view. Everything I have described is something one insurer can do alone. But imagine the whole market measuring claims quality the same way, against one shared line. The industry would have, for the first time, a common language for the thing it has always found hardest to talk about. Not each firm marking its own homework, but a single, external yardstick that a customer, a leader, an investor, and in time perhaps a regulator, could all read the same way. The investor belongs in that list for a hard-headed reason. Once

claims quality is measured, it becomes a leading indicator of retention economics, of premium that will stay and premium that will have to be expensively re-bought, and that is a language capital markets already read fluently.

That would matter for a reason bigger than any one company's advantage. Trust in insurance is, to a degree, shared. When one insurer fails a claimant badly, the whole industry wears a little of the blame, because the customer rarely distinguishes. So making claims quality visible, and comparable, and therefore rewarded, would lift the floor for everyone. It would turn good claims handling from a slogan no one believes into a fact that can be seen and chosen. In a market where machines are increasingly optimising cost behind the scenes, a shared, human measure of trust is not a nicety. It may be the thing that keeps the industry honest.

I believe that is where this should go, and I will not pretend to be a disinterested party, because I have spent years building towards exactly this kind of measure, and you are entitled to weigh that. But the case does not rest on me. It rests on a simple choice facing the industry. Either claims quality stays invisible, and the machines optimise cost in the dark, or the industry decides to see it, to share a language for it, and to compete on it in the open. The second is the better world, for customers and for the firms that serve them. And a shared standard only ever emerges because leaders decide they want it. So it is, in the end, a leadership question, like everything else in this book.

The next chapter of this industry will be written by leaders who use the machine to elevate the human rather than erase them, and who are confident enough in their claims to be measured against everyone else. That takes nerve. It takes belief in the very thing this whole book has been about. And before I say the last thing I have to say, the yardstick's two newest readers deserve chapters of their own. Because the regulator, for all its data, is currently missing the part of the story that matters most, and the investor, for all their models, has never been shown the number

that predicts an insurer's future book. Let me take them in turn.

Chapter 14: What the regulator cannot see

Imagine a careful driver. She has owned the car from new, serviced it on schedule, kept it garaged, resisted the kerbs that claim everyone else's alloys. Then one morning someone runs into it, and it is written off. Her insurer accepts the claim promptly, values the car, deducts a standard amount for assumed pre-existing damage, and pays. She takes the settlement, because the insurer is the expert, and discovers it will not buy her an equivalent car. What she never learns is that the deduction was automatic. Nobody inspected anything. The system simply assumed her car carried the wear of an average car, and priced the assumption into her payout. The more carefully she had looked after it, the more the assumption took from her.

Now multiply her. In September 2025, the regulator announced that insurers would pay around £200 million in compensation to some 270,000 motorists for exactly this practice: automatic deductions for assumed pre-existing damage on written-off and stolen vehicles. Eighteen firms were involved, covering roughly ninety per cent of the market. By the time of the announcement, £129 million had already been repaid to nearly 150,000 customers. This was not a rogue firm. It was close to standard practice, across nearly the whole market, and it ran until the regulator dug it out through a multi-firm review, having

first warned the industry about undervaluation nearly three years earlier.

I tell this story at the start of the regulator's chapter for one reason, and it is not to bash the industry; chapter 12 earned better than that. I tell it because of where the practice hid. Think about every instrument the industry and its regulator were watching while this was happening, and ask which of them could have caught it.

Not acceptance. Every one of those claims was accepted. Paid, promptly, and counted as a success on the very measure the industry treats as its headline reassurance. Acceptance answered its own question perfectly: did the firm pay? Yes. It cannot see what the payment should have been. A claim can be accepted and still be unfair, and here were a quarter of a million of them, all sitting on the right side of the acceptance rate.

Not average payouts. An automatic deduction shaves every settlement by a little. The average moves by pennies against a backdrop of car values that swing with the used-car market. No supervisor scanning a firm's average payout would see a quiet, systematic subtraction from careful drivers; the signal drowns in the noise of the market itself.

Not complaints. The whole mechanism depended on the customer not knowing. The careful driver had no idea an assumption had been made about her car, so she had nothing to complain about. We saw in chapter 8 that complaints measure the willingness to fight; you cannot fight what you cannot see. The claimants most wronged by this practice were, by construction, the least likely to complain about it.

Every output dial read normal. The unfairness was systemic, sizeable, and invisible to all of them, and it stayed invisible until the regulator went looking the hard way. That is not a story about one bad practice. It is a story about a measurement system, and it is why this chapter exists.

Let me be precise and fair about what the regulator can see, because its instruments are valuable and this book has already leaned on them. Since 2021 the FCA has required firms to report, and has published, its Value Measures: how often customers claim, what share of claims are accepted,

what the average payout is, and how many complaints arrive per claim. Four dials, per firm, in public. Chapter 2 built its argument on them, and the daylight they bring is real. But look at the grammar of those four dials. Every one of them records what the firm did. Frequency: how often it was asked. Acceptance: whether it paid. Payout: how much. Complaints: how many customers fought back. There is no dial for what the customer got, in the sense that a person means it: whether the settlement rebuilt their life, whether the process left them trusting the promise, whether they would stake next year's premium on it again. The regulator, with the best output data in the world, is reading the firm's account of the transaction. The customer's account of the relationship is not on the dashboard.

And the regulator, to its credit, says as much where its data is weakest. Its own publication of the 2024 figures cautions that home acceptance rates ranged from around half of claims at some firms to nearly all of them at others, and that these comparisons "may not be like for like", partly because firms interpret the reporting rules differently, including how a declined claim is even recorded. Sit with that for a moment. On the regulator's headline dial, the industry cannot fully agree what a declined claim is. The instrument is honest about its own blur.

Now set beside those four dials what this book has been reading. Not what firms did, but what customers report: how the claim felt, stage by stage, and, because the tracker follows them to renewal, what they did next, at the same price. Call the first kind of data outputs and the second kind outcomes, because that is the regulator's own language. Since the Consumer Duty arrived, firms are no longer judged only on process compliance; they are required to deliver good outcomes, and the regulator has said it will act faster and more firmly where outcomes fail. Its current priorities say the quiet part out loud: claims must be handled promptly, fairly and transparently, and deliver good outcomes, with claims handling investigations under way this very year. Here is the awkwardness in that position, stated respectfully. The Duty demands outcomes. The instruments report outputs. A firm can

evidence every output dial and still not know, and not be able to show, how the claim landed on the person, or what the person did about it. Consumer-reported outcome data is not a rival to the regulator's ledger. It is the missing column: the two read together answer the question neither can answer alone, whether the firm that pays claims also keeps its promises in the way the customer actually experiences them.

Would a consumer-outcome line have caught the total-loss problem sooner? I want to answer that carefully, because this book does not deal in hindsight heroics. Honestly: not directly, and the appendix says why. The survey does not yet record how a claim was settled, so it could not have isolated total-loss claimants as a group; that is one of the survey changes we have asked for, and the closing pages return to it. But indirect tripwires are a different matter. A systematic underpayment lands somewhere in the data as bruise even when its cause is invisible: in satisfaction that sags for one firm's claimants while its outputs gleam, in customers reporting that the outcome was not what they expected, in the quiet drift away at renewal that no complaint ever explained. An output dashboard has no tripwires at all for a practice the customer cannot see. An outcome line has at least the bruise. Regulation by outputs alone means finding these things the way this one was found: years later, through a multi-firm review, £200 million after the fact. The cheaper discovery mechanism is listening, at scale, to what claimants say and watching what they do.

Which brings me to the question I have owed the regulator since chapter 6, and it is the question only this kind of data can pose, let alone answer. Four in ten customers stay after a poor claim. Stay, after the promise was tested and, by their own account, broken. Chapter 6 offered two readings: forgiveness, or inertia. Put on the regulator's spectacles and feel how much turns on which it is. If the stayers are contented forgivers, weighing one bad episode against years of good service, there is no issue; that is a market working. But if they are staying because leaving takes energy, confidence and comparison-site fluency they do not have, then something regulatory is happening: the ordinary market

discipline that should punish poor claims handling is being absorbed by the customers least equipped to impose it. Poor handling would face its consequences only from the customers who can afford to deliver them, and be subsidised by those who cannot. Retention, the number every board celebrates, would be functioning, for this group, as a measure of resignation.

The survey carries the markers to test this, and we have now run the test. On every switching-capability marker that shows a real signal, the poor-claim stayers are not the trapped. They use comparison sites more than the average customer who stays, not less. They shopped around at the very renewal where they chose to stay, at a higher rate than never-claimed stayers did. They are no older than anyone else. Whatever is keeping these customers with an insurer that handled their claim badly, it is not an inability to work the market; they engage with it at least as much as everyone else, and they stay anyway. So the uncomfortable reading, market discipline absorbed by those least equipped to impose it, does not hold on the measures available here, and I am glad to be able to report that rather than merely pose it. Two caveats belong beside the finding. The data describes who stayed, not why; no causal story is proven, and I will not print one. And the markers the survey can see are behavioural, comparison-site use, shopping activity, age, not a full vulnerability assessment; the question deserves to be asked again as the data deepens. But the first pass at precisely the question the Consumer Duty was written to ask has come back, and it has come back in the market's favour. What I want you to notice is not just the answer. It is that the question was answerable at all. Who stays, who leaves, and whether staying was a choice: that is Consumer Duty territory, and it lives in outcome data or nowhere. No output measure could even have posed it.

I said in chapter 13 that a shared yardstick could one day be read by a customer, a leader, an investor, and in time perhaps a regulator. This chapter is my case for the last reader on that list. Not that the regulator should abandon its ledger; it should not. But the next total-loss scandal,

whatever it turns out to be, is sitting right now in some gap between what firms report and what customers experience. Only one kind of instrument watches that gap. The regulator has more data than it has ever had. What it cannot yet see is the part of the story this book has been telling: how the promise felt, and what the customer did next.

There remains one reader from chapter 13's sentence, and on the surface he is the regulator's opposite: he does not care how the claim felt, only what the book of business is worth. I am going to argue that the same number serves him too, and that claims quality may be the best signal of an insurer's future earnings that the market has never priced. To the investor, then.

Chapter 15: Reading claims like an investor

Of all the readers I imagined while writing this book, the one I am about to address is the one who would least expect to be here. The investor. The analyst building a model of an insurer, the fund manager weighing one carrier against another, the buyer running diligence on a book of business. You do not care, professionally speaking, how a claim felt. You care what the company is worth. Stay with me anyway, because I am going to show you a number that predicts something you are currently forced to guess, and it is sitting outside the accounts you read.

Start with what you can see. An insurer's statutory reporting will tell you almost everything about claims as a cost. Loss ratios, combined ratios, reserve development, large-loss experience, inflation assumptions: the whole apparatus of claims-as-outflow, disclosed, audited and comparable. What none of it tells you is what those claims did to the book. Here is the thought experiment that opens the gap. Take two insurers with identical loss ratios. Both paid out the same share of premium in claims last year; on the cost side of the ledger they are twins. But one of them handled those claims well, and the other handled them badly. Everything in this book says their futures are different. The first insurer's claimants renew; across two years, roughly seven in ten of its good-claim motor customers stayed. The second insurer's claimants

leave; nearly six in ten walked after a poor claim. Same loss ratio. Opposite trajectories. And nothing, anywhere in either firm's published numbers, lets you tell them apart. The cost of every claim is disclosed to the penny. The consequence of every claim is invisible.

Why should that consequence matter to a valuation? Because of what retained premium is. Retention is the cheapest growth there is; the customer has already been won, the acquisition cost already spent, and every renewal arrives without a pound of new marketing. A book that keeps its customers compounds. A book that loses them must rebuy its own revenue every year, paying acquisition costs on customers it used to own, running hard to stand still. Investors in other industries price this distinction obsessively. Ask anyone who values software businesses: two companies with identical revenue are given entirely different multiples depending on how much of that revenue returns by itself. Recurring revenue is worth more than re-bought revenue, per pound, every time, because it is cheaper to keep and more predictable to hold. Insurance renewal books are exactly this kind of asset, an annuity stream of premium, and yet the market prices insurers with barely any visibility of the single biggest behavioural driver of whether the annuity persists. This book has spent fourteen chapters identifying that driver. It is the claim.

Let me put the dividend in your language, using the arithmetic of chapter 7 and its stated, illustrative assumptions. Take a motor book of one million policies. The difference between handling its claims well and handling them badly is worth in the region of £10.6 million of retained premium a year, and about £1.5 million of profit. Read those numbers the way you would read any earnings-quality adjustment. They are premium the good-claims book keeps that the poor-claims book must go out and rebuy, at full acquisition cost, every single year. They recur. They compound with tenure. And, crucially for your purposes, the gap driving them survives a price control, so it is not a soft satisfaction artefact that dissolves when the pricing team gets aggressive; it is a behavioural difference between books that persists when price is held

like for like. If you knew which of two otherwise identical insurers sat on which side of that gap, you would not value them the same. Today, you do not know. That is the pitch, and I mean pitch in the analyst's sense: an information edge, lying unexploited.

Now for the property that should interest you most: timing. Everything in statutory reporting is a trailing indicator. This year's poor claims handling shows up as next year's weaker retention, which shows up as the year after's higher acquisition spend and slower growth, which an analyst finally notices as margin pressure two annual reports later. The chain takes years to surface, and by the time it does, the cause is unrecoverable and unattributed. Claims experience data runs at the other end of the chain. It is collected while the claims are happening, from the claimants themselves, and it predicts what they do at the next renewal. It is, in the strictest sense, a leading indicator of retention economics, visible one to two years ahead of anything the company will file. And because the tracker measures every firm the same way, from outside, it is comparable across the market in a way no company's internal survey could ever be. A market-wide line, updated continuously, on which every major insurer's claims quality can be read like for like: that is not a customer-experience curiosity. That is an input a model can take, a screen a fund can run, a trend a holder can track between reporting seasons.

If you are running diligence rather than running a screen, the questions almost write themselves. Where does the target sit on the market line, and which way has it been moving? What share of its claimants come away neutral, the quiet middle of chapter 6, since indifference leaks value that no complaint log will disclose? How does its brand experience compare with the experience delivered by the claims machine behind it, which chapter 10 showed can be two different things? And has it been buying its combined ratio with the customer's patience? That last question deserves its own paragraph, because it is where an investor's incentives and this book's argument collide.

There is a way to make claims cheaper that works every time, for

one year. Squeeze the experience: stretch the timescales, trim the courtesies, automate the human moments, tighten every settlement. The loss ratio improves immediately. The damage, as this book has priced at length, arrives at the next renewal and the one after, in a book that quietly empties. An insurer doing this is not improving its earnings; it is borrowing from its own future book at a ruinous interest rate, and the loan is invisible in the accounts precisely because the collateral, customer trust, appears nowhere on the balance sheet. To a short-term holder, that squeeze looks like discipline. To anyone holding longer than a renewal cycle, it is the classic earnings-quality trap: current margin flattered by the consumption of an unbooked asset. The claims-quality line is the only instrument that catches it in the act, because it shows the experience deteriorating while the reported numbers still gleam. If you take one habit from this chapter, take this: whenever a personal-lines insurer reports a striking improvement in claims costs, ask what happened to its claimants' experience in the same period. If the answer is "it fell", you have not found efficiency. You have found the future book, being spent.

The same lens sharpens the biggest transactions this market makes. Personal lines has been consolidating; large books of customers now change hands with their claims operations attached. Every one of those deals prices reserves, renewal rights, brand value and synergies. Almost none of them can price what the target's last three years of claims handling did to the loyalty of the book being bought. A buyer inherits banked trust or banked resentment, millions of customers who have either had the promise proven or had it broken, and that inheritance starts paying out, or leaking, from the first renewal after completion. It is knowable, from data like this book's, and it is material, by chapter 7's arithmetic. Diligence that stops at the loss ratio is buying the machine without checking what the machine has been doing to the customers.

I should be clear about what I am not claiming. I am not claiming claims quality is the only driver of retention; price, brand and distribution all move the number, which is exactly why the dividend was measured

with price held constant. I am not claiming the market line replaces the accounts; it complements them, the way any leading indicator complements a trailing one. And the pound figures here carry every assumption and caveat of chapter 7, stated there and in the appendix, because an investor of all readers should refuse a number whose workings are hidden. What I am claiming is narrower and, I think, harder to dismiss: there exists a measurable, comparable, price-controlled behavioural signal that leads retention economics by a year or more, and the capital markets that price these companies have never had it. In chapter 13 I argued the industry should want a shared yardstick for claims quality because it would keep the machine age honest. Here is the same argument from your side of the table: what keeps firms honest is what gets measured and priced. The day claims quality is visible to capital, every board in the industry will discover a sudden, sincere interest in the moments this book has been describing. Customers will be better served because you, of all people, started watching. That is not a paradox. That is how markets are supposed to work.

And with that, the cast is complete. A customer, a leader, an investor, and in time perhaps a regulator, all reading the same line, each for their own reasons, each keeping the others honest. Four readers of one number, and the number measures a promise. Which returns us, at last, to where we began. I opened this book with a confession about promises, and I owe you an ending that keeps one. Let me gather everything we have seen, and say what it amounts to.

Chapter 16: The promise kept

I began this book with a confession. That I have spent my career building promises, and that the one place a promise is truly tested is the one place I could never reach. The claim. We have spent the chapters since travelling into that moment together, and I want to end where we started, because the journey has, I hope, changed what the word means.

We have seen that the claim is the only test that truly counts. That almost everyone passes its lowest bar, the paying, and that the real contest lies in how. That the distance between a good claim and a poor one is somewhere between twenty-six and forty-plus points of a customer's future, at the same price. That trust is banked at the claim and spent at the renewal that would otherwise have lost the customer, and that home is its own harder story, where even a good claim leaves a bruise and the prize is halving the damage. That fine is not fine, and the quiet middle leaks more loyalty than the angry tail. That the whole dividend can be priced on a single sheet of paper, with your own numbers, and that the tools the industry trusted were never built to see any of it. We have seen that the moments which decide a claim are human moments, small and easily missed. That leading them well means holding process and people together, under tension, at once, inside an industry squeezed hard by rising costs and industrial fraud. That the machine now arriving does

not end the human's part in this, but makes it matter more. And we have seen the claim through two new pairs of eyes: the regulator's, which cannot yet see the part of the story that matters most, and the investor's, which has never been shown the number that predicts the future of the book.

Underneath all of it has run a single idea, borrowed, as you will remember, from the authors of The Trusted Advisor, and pressed into the service of claims. That trust is credibility, and reliability, and intimacy, divided by self-orientation. Four plain questions a customer asks, silently, at the worst moment of their year. Do you know what you are doing. Will you do what you said. Am I safe with you. And, beneath them all, when it comes to it, are you on my side, or your own. A claim is simply the moment those questions are answered, in full, whether the customer meant to ask them or not.

So the real subject of this book was never claims, and never data, and never machines. It was that last question. Whether, at the moment it mattered most, your customer decided you were for them. Everything else, the measurement, the journey, the leadership, the technology, is only ever in service of getting that one answer right, at scale, again and again, on the days your customers will remember for the rest of their lives.

Think back to Mrs Miggins, whose claim went perfectly and who was not happy at all. She is the whole book in one small story. Everything that could be counted was done, and the thing that mattered was missed, because no one asked what she needed. The gap between a good claim and a poor one is very often exactly that small, and exactly that human. The gap is rarely a matter of money or systems. It comes down to whether someone was paying attention to the person.

I have tried throughout to be a champion rather than a critic, and I will end that way too. The claim has been treated, for a long time, as a cost to be contained. I hope I have persuaded you that it is something else entirely. It is the single greatest opportunity most insurers have to build the trust that keeps customers, earns growth, and cannot be copied by

a rival with a bigger budget. That opportunity has been sitting in plain sight, unmeasured and unclaimed. It is there for the leaders who choose to see it.

Which brings me, finally, to the sentence this book was named for. For years I thought the promise was the thing we made at the sale, in the advertising, in the brand. I was wrong. The promise is not made at the sale. It is kept, or broken, at the claim. The claim is the promise. Keep it well, and everything an insurer wants, loyalty, growth, trust, a reputation a competitor cannot buy, follows from it.

That is the test that counts. It always was. I hope you will go and win it.

Appendix A: A note on the numbers

A book that asks a board to act on its figures owes that board the figures in full: where each one comes from, how many people stand behind it, how wide the uncertainty runs, and what it does and does not mean. That is what this appendix is for. Nothing here is needed to follow the argument. Everything here is needed to check it.

Unless a line says otherwise, every survey figure in this book is drawn from the Insurance Behaviour Tracker, motor and home read separately, over the window used for the dividend analysis, and reported unweighted. Confidence intervals are ninety-five per cent. Where a base is small, the figure is flagged, and where it is very small it is not quoted at all. Appendix B explains the method behind all of this; this appendix is the ledger.

The words behind the numbers

- **Good, neutral, poor.** A claimant's own satisfaction with the claim, on a five-point scale. Good is the top two boxes, four or five. Neutral is the middle, three. Poor is the bottom two, one or two.

- **Never-claimed.** A customer who reports no claim in the relevant window. The natural comparator for what a claim does.
- **Switched, and shopped.** Switching is leaving the insurer at renewal. Shopping is looking around, whether or not the customer then left. They are different measures and are never treated as the same.
- **The two reads.** A twelve-month read follows customers over one renewal; a two-year read follows them over two. The book keeps them apart, because they answer different questions, and it never describes one as the other.
- **Price bands.** The customer's own recollection of how their renewal price changed. Recollection, not a ledger; that caveat travels with every price-controlled figure.
- **Base tiers.** A cell of fifty respondents or more is treated as firm. Between thirty and fifty it is shown but flagged as indicative. Below thirty it is not quoted.

The core dividend

Figure	Value	95% interval	Base	Read
Motor, good-claim switch rate	29.6%	27.2–32.1	1,342	2-year
Motor, poor-claim switch rate	58.2%	48.3–67.4	98	2-year
Motor, good-vs-poor gap	28.6pp ("twenty-nine points")	—	—	2-year

Figure	Value	95% interval	Base	Read
Home, good-claim switch rate	33.9%	31.4–36.4	1,393	2-year
Home, poor-claim switch rate	60.0%	49.4–69.8	85	2-year
Home, good-vs-poor gap	26.1pp (“twenty-six points”)	—	—	2-year
Motor, never-claimed switch rate	30.1%	28.7–31.6	4,027	2-year
Home, never-claimed switch rate	23.0%	21.8–24.2	4,553	2-year
Price-controlled gap, motor	~30pp	—	—	AME
Price-controlled gap, home	~25pp	—	—	AME
Satisfied-vs-dissatisfied renewal gap (motor)	42.8pp (“beyond forty points”)	+33.1 to +52.1	109 / 1,827	linkage cohort

The price-controlled gap holds price change, age band and region like for like, and reports the average marginal effect. It is the number the book

leans on, because it is the one that survives the obvious objection that loyalty is really pricing in disguise. The satisfied-vs-dissatisfied gap is a wider net over a different cohort; the book treats it as the outer bound and the price-controlled gap as the floor, and never adds the two or swaps one for the other.

One caution worth stating up front, because it is the easiest mistake to make in good faith. There is more than one defensible way to measure "how much more poor claims lose than good ones," each resting on a different cohort. The two-year figures above are the book's spine. The satisfied-vs-dissatisfied figure is a separate, wider measure. Wherever the book quotes a gap, it names which one it means.

The fine middle, and who stays

Figure	Value	95% interval	Base
Motor, neutral-claim switch rate (2-year)	55.9%	48.4–63.1	170
Home, neutral-claim switch rate (2-year)	49.7%	41.9–57.5	153
Neutral-to-dissatisfied ratio, motor	1.85 : 1	—	703 / 379
Neutral-to-dissatisfied ratio, home	2.25 : 1	—	660 / 293
Share of excess churn from the middle, motor	63%	—	—

Figure	Value	95% interval	Base
Share of excess churn from the middle, home	57%	—	—
Silent detractors: poor-claim stayers who are detractors, motor	70.7%	55.5–82.4	41 (indicative)
Silent detractors: poor-claim stayers who are detractors, home	64.7%	47.9–78.5	34 (indicative)
Good-claim stayers who are detractors, motor / home	5.0% / 3.7%	—	945 / 921

The motor and home neutral rates are deliberately described differently in the body. In motor the neutral rate and the poor rate are statistically indistinguishable, and the book says so. In home the neutral rate sits clearly below the poor rate, and the book keeps its more cautious "roughly half, against six in ten" rather than forcing the two into equality. The silent-detractor cells are indicative, not firm; the book flags this wherever the figure appears.

Advocacy, attribution and the journey

Figure	Value	Note
Net recommendation, good vs poor, motor	+55 vs −25	swing ~80 points
Net recommendation, good vs poor, home	+61 vs −48	swing ~109 points
Claims-driven shopping reasons cited, motor: claimants vs never-claimed	~32–36% vs 17.8%	any claim roughly doubles citation
Claims-driven shopping reasons cited, home: claimants vs never-claimed	~40–47% vs 15.3%	any claim roughly triples citation
“Cover was not what I expected” (home), claimants vs never-claimed	22–25% vs 9.4%	the underinsurance code
Journey statement most separating stayers from leavers, motor	“knew what to expect”, +27.0pp	20.9–33.4
Journey statement most separating stayers from leavers, home	“easy to make a claim”, +18.3pp	12.7–24.5
Reliability vs intimacy, head-to-head	motor +2.9pp, home +0.3pp	both intervals include zero: no winner

Two of these deserve a word. On the shopping reasons, the finding is that having claimed at all drives citation; a poor claim does not drive

it measurably more than a good one. The book says exactly that, and no more. On reliability versus intimacy, the head-to-head is a statistical dead heat in both products; the book reports that the human half of a claim carries as much of a customer's future as the procedural half, and crowns neither.

The whole market, and the blind spots

Figure	Value	Note
Market claimant satisfaction, motor / home	81.4% / 83.0%	top-two-box
Best-to-worst firm satisfaction spread (market-wide)	over 20 index points	the "some twenty points" of chapter 2's argument
Acceptance-vs-satisfaction spread, matched FCA firms (motor)	18.6 points across 11 underwriters	the Chapter 2 contrast; motor only
Claims handled by another insurer (the hidden tenth), motor	10.7% (~1 in 10)	home 4.3%
Tracker annual claim rate, motor / home	6.2% / 6.0%	the rate behind chapter 7's arithmetic

The market-wide satisfaction spread and the matched-firm spread are two different measures on two different populations. They land near the same number by coincidence, not by identity, and the book keeps them distinct. The Chapter 2 contrast is drawn from motor alone, because in home the regulator's own acceptance measure really does vary from firm to firm, and the flat line it depends on is not there.

The decay curve

The good-vs-poor gap, cut by years since the claim, both products, is quotable through year three and fades into statistical noise by year four:

Years since claim	Motor gap	Home gap
1	37.6pp	32.6pp
2	22.9pp	21.8pp
3	19.6pp	17.9pp
4	7.6pp (not significant)	8.1pp (not significant)

The pound figures

The book's one deliberate silence is the pound figures. Chapters 7 and 15 turn the loyalty gap into money, and the method they use is set out there in full, line by line, so that any reader can rebuild the sums with their own premiums, margins, acquisition costs and claim rates. The specific worked figures I could have printed rest on planning assumptions of mine, not measurements, and until those assumptions have cleared internal sign-off they are held back rather than printed as though they were facts. This is the same discipline the rest of this appendix follows: show the reader how the number is made, and never hand over a figure whose workings are hidden. The gap itself, the one input that is measured rather than assumed, is in the tables above. The rest is arithmetic you should do with your own book.

The standing caveats

Every figure in this book is offered subject to five standing limits, and they are worth carrying in mind together.

The data is unweighted, consistent with the analysis pack it is drawn from. Perceived price change is the customer's recollection, not a

billing record, so every price-controlled figure inherits that softness. Most figures describe one renewal cycle; only the decay curve looks across several. The tracker reaches a broader spread of claimants than a post-claim survey, which is its central virtue, but no survey removes self-selection entirely. And some cuts rest on partial coverage: the claim-satisfaction question is answered by roughly half of claimants, and the insurer who handled a claim is identifiable for around eight in ten. None of these overturns a finding in this book. All of them are the reason the book states its numbers as ranges, flags its small bases, and prefers the conservative reading wherever two are available.

Appendix B: The method

The argument of this book rests on one instrument, and an argument is only as trustworthy as the instrument beneath it. So let me describe how the evidence is gathered, why it can see things other instruments cannot, and where its own limits lie.

The tracker

The Insurance Behaviour Tracker is a continuous study of ordinary insurance customers, running for more than six years, covering motor and home. Over that time it has heard from more than 150,000 customers, over 30,000 of them claimants describing a real claim they had recently lived through. It is not a one-off poll. It runs month after month, which is what lets it follow customers forward and watch what they actually do.

Why it is not a claims survey, and why that matters

This is the feature that does the most work in this book, so it is worth being precise about. Most instruments that measure claims are triggered by a claim. A complaints process hears from the aggrieved. A review site hears from the delighted and the furious. Even a well-run post-claim

survey is sent by the insurer being judged, and answered most eagerly by the people who felt something strong. Each of these is tuned to intensity, and each is therefore deaf to the same people: the great, quiet middle whose claim was simply fine, who have no story to tell and no reason to tell it.

The tracker is built the other way round. It asks customers about their insurance in the round and comes to their claims experience along the way, with no claim as the reason for the conversation. That single design choice is why it can hear the middle, and the middle, as chapter 6 shows, is where a surprising amount of loyalty is quietly won and lost. It is also why the tracker can be laid against behaviour at renewal without the circularity that dogs a claim-triggered survey.

Following customers to renewal

Because the study is continuous, a customer's account of a claim can be tied to what they did at their next renewal: stayed, or switched. This is the join on which the whole dividend rests. It is not what a customer said they might do; it is what they did. And because the study also captures the customer's recollection of how their renewal price moved, the loyalty effect can be read like for like among customers who saw the same price change.

Holding price still

The obvious objection to the whole argument is that loyalty is really just pricing wearing a costume: perhaps the customers who left a poor claim were just the ones whose price rose. The method answers this by comparing like with like. Customers are grouped by their perceived price-change band, by age band and by region, and the claims effect is read within those groups, reported as an average marginal effect. The effect survives, and barely narrows. The one caveat that travels with

every such figure is that the price change is the customer's recollection, not a billing record; it is a proxy, and an honest one, but a proxy.

The matched regulatory extract

For two of its arguments the book joins the tracker to the regulator's own published data, the FCA Value Measures: for each firm, how often customers claim, what share of claims are accepted, the average payout, and complaints per claim. The two datasets are matched firm by firm. Because one authorised firm often underwrites several consumer brands, the match is made at the level of the underwriter as well as the brand, and the resulting line can be read both ways: by brand, to see what customers chose and felt, and by underwriter, to see what one claims operation delivers across everything it touches. The matched set is necessarily smaller than either source alone, since a firm has to appear in both with enough claimant volume to give a reliable reading; at underwriter level it runs to roughly a dozen firms per product. Where the book quotes this extract, it says so.

Suppression, intervals and verification

Three rules keep the numbers honest. First, every proportion is reported with a ninety-five per cent confidence interval, and every comparison between groups with an interval on the difference. Second, base sizes are tiered: a cell of fifty or more is treated as firm, a cell between thirty and fifty is shown but flagged as indicative, and a cell below thirty is not quoted. Third, every new figure is reconciled against the already-verified figures before it is used, so that a fresh cut cannot silently drift from the numbers the book already stands on. Anything that cannot be verified is left flagged, kept qualitative, or posed as a question, rather than dressed up as a finding.

What the instrument cannot do

Honesty about a method means naming its blind spots. The claim-satisfaction question is answered by roughly half of claimants, and the insurer who handled a claim is identifiable for around eight in ten, so some cuts rest on partial coverage. Claim timing resolves to whole years, not months, which is why the decay curve of chapter 7 is read year by year and cannot yet locate a within-year window. The data does not record how a claim was settled, so total-loss claimants cannot yet be isolated as a group. And the study covers motor and home; the claims questions do not extend to other products. These are not reasons to distrust the findings. They are the edges of what the instrument can currently see, and appendix C is the list of the ones we are working to push back.

Appendix C: Open questions

A study that pretends to have answered everything has stopped being a study. This book has been able to settle questions that were open when I began writing it, the half-life of the dividend among them. But it has opened others in the asking, and a few important ones it cannot yet answer at all. I would rather set them down plainly than leave the impression that the account is closed. Each is a piece of work either underway or asked for, and each is a reason there will be more to say.

The total-loss question

The book shows, in chapter 14, how a systematic underpayment on written-off cars stayed invisible to every output measure the industry and its regulator watch. The tracker would have carried at least a bruise of it. But it cannot yet isolate total-loss claimants as a group, because the survey does not record how a claim was settled: total loss, repair or cash. We have asked for a settlement-type field to be added to a future wave. Until it is confirmed and fielded, the total-loss retention question, whether writing a car off costs an insurer more loyalty than repairing it, remains one the data cannot directly answer.

The intervention window

Chapter 7 shows the dividend decaying over three to four years, read year by year. What it cannot yet show is the shape within the first year: the weeks and months just after a claim, when the banked trust is at its freshest and an insurer's chance to convert it into longer loyalty is greatest. Claim timing in the current data resolves to whole years, so the within-year window is invisible. A finer, months-since-claim measure is the second change we have asked for. With it, the "act inside the window" instruction of chapter 7 could be sharpened from a principle into a timetable.

Beyond the fourth year

The decay curve runs out at year four, not because the dividend necessarily ends there, but because the survey's memory does: there are too few claims that far back in the current window to read the gap cleanly. As the study accumulates further waves, the curve can be extended, and the question of whether a trace of the dividend persists at five years and beyond becomes answerable. That is second-edition territory, and it is named as such in the body.

The household halo

A household rarely holds one policy. The car, the home, sometimes more, often sit with the same insurer or move together. If a good claim on one policy lifts loyalty on the others, or a poor one drags them down, the dividend this book has priced on a single policy is an understatement. The tracker holds enough to begin testing this; it is a natural next cut, and a commercially significant one, because it would change the size of the prize.

Beyond motor and home

The claims argument in this book is built on motor and home, the two products where the tracker's claims evidence is deepest. Whether the same dividend, the same quiet middle, the same decay, appears in travel, in pet, in the commercial lines, is an open and answerable question, and one the method transfers to directly. The principle, that the claim is the moment the promise is tested, is not specific to a car or a house. Whether the numbers hold the same shape elsewhere is worth finding out.

None of these gaps weakens the argument the book has made. They are the frontier of it. A living measure of claims quality is not a monument to be unveiled once and admired; it is an instrument that gets sharper every year it runs, and answers next year questions it could only pose this year. That, in the end, is the case for treating claims experience as something to be measured continuously rather than surveyed occasionally. The questions do not run out. Neither should the listening.